CONFLUENCE
DOCUMENTATION

Administrators Guide

Confluence 3.0 Administrator's Guide
Configuring Confluence Data and Backups System Administration Importing Data Mail Configuration Security User Management Design and Layout Integrating Confluence and JIRA Plugins and Macros Performance Tuning Character Encoding Support

Additional Resources
Visit the Configuration Guide for documentation on configuring databases and application servers. The Confluence User Guide has information on how to use Confluence as a collaborative tool. Go to Documentation Home for links to more resources.

Download
You can download the Confluence Admin Guide in PDF, HTML or XML formats.

Site Administrator?
The **Confluence Administrators Guide** provides information to site administrators on how to manage their Confluence instances. If you still have a question that hasn't been answered, write and tell us about it. If you are using Confluence Hosted for Small Business, please refer to the Administrator's Guide to Confluence Hosted for Small Business.

Configuring Confluence

Site Configuration
Configuring the Site Homepage
Editing the Site Title
Editing the Site Welcome Message
View Space Goes to Browse Space
Editing the Global Logo
Configuring the Server Base URL
Configuring the Site Support Address
Configuring HTTP Timeout Settings
Configuring System Properties
Customising Default Space Content

Optional Settings
Enabling Remote APIs
Enabling Rich Text Editing Option
Making Rich Text Editing default
Enabling CamelCase Linking
Enabling Trackback
Enabling Threaded Comments
Attachment Storage Configuration
Configuring Attachment Size

Display Settings
Configuring Indexing Language

Data and Backups

System Administration

Importing Data

Mail Configuration

Security

Hiding External Links From Search Engines
Excluding external referrers
User Email Visibility
Anonymous Access to Remote API
Adding SSL for Secure Logins and Page Security

User Management

Confluence User Management
Global Groups Overview
Global Permissions Overview
Setting up Anonymous Access
Adding a New User
Editing User Details
Removing a User
Enabling or Disabling Public Signup
Adding or Removing Users in Groups
Adding a Group
Removing a Group
Viewing members of a group
Restoring Passwords To Recover Admin User Rights

External User Management
Understanding User Management in Confluence
Disabling the Built-In User Management
LDAP Integration Overview
Troubleshooting LDAP User Management
Requesting External User Management Support
Troubleshooting the "Not Permitted" Screen under LDAP Integration

Crowd User Management
Integrating Confluence with Crowd

JIRA User Management
Delegating User Management in Confluence to JIRA
Override properties in JIRA to Confluence Bridge

Design and Layout

Configuring Layouts
Customising Look and Feel Overview
Customising Colour Schemes
Customising Layouts
Working With Decorator Macros
Global Templates
Customising the eMail Templates

Configuring Themes
Themes Overview
Creating a Theme
Applying a Theme To A Site
Including Cascading Stylesheets in Themes

Integrating Confluence and JIRA

Integrating JIRA and Confluence
Configuring Jira Issues Icon mappings
Adding Confluence EAR-WAR to JIRA Standalone
Setting Up Trusted Communication between JIRA and Confluence

Plugin Management

Confluence Plugin Guide
Installing and Configuring Plugins Manually
Installing and Configuring Plugins using the Plugin Repository Client
Configuring the Office Connector
Setting Up Trusted Communication between JIRA and Confluence
Configuring a URL Whitelist
Enabling the 'Mail Page' plugin

Macros
Configuring User Macros
Editing and Removing User Macros
Enabling HTML macros
Enabling the html-include Macro

Flowchart Macro
Troubleshooting the Gallery Macro
Configuring the userlister Macro
Working With Decorator Macros
Writing Macros

Performance Tuning

Memory usage and requirements
Configuring a Large Confluence Installation
Performance tuning
Enabling HTTP Compression
Working with Confluence Logs
Page Request Profiling
Profiling using the YourKit Plugin

Character Encoding

Configuring Encoding
Troubleshooting Character Encodings

Support

How to Get Support

Administration

- Cache Statistics
- Changing time of Daily Backup
- Confluence Data Directory Configuration
- Confluence home directory contents
- Content Index Administration
- Finding Unused Spaces
- Important Directories and Files
    - Confluence Home Directory
    - Confluence Installation Directory
- Manually Backing Up The Site
    - Configuring Daily Backups
        - User Submitted Backup & Restore Scripts
- Migrating Confluence Between Servers
- Rebuilding the Ancestor Table
- Restoring a Site
- Restoring a Space
    - Changing the version of a space backup
- Restoring a Test Instance from Production
- Restoring Data from other Backups
- Restoring Data from the Administration Console
- Retrieve file attachments from a backup
- Troubleshooting failed XML site backups
- Troubleshooting XML backups that fail on restore
    - Migrating from HSQLDB to MySQL
- Viewing and Editing License Details
- Viewing System Information
    - Live Monitoring Using the JMX Interface
    - Site Statistics
- Installing Patched Class Files

Administrators Guide Home Confluence Documentation Home

Cache Statistics

Confluence provides statistics about its internal caches that allow you to track the size and hit ratio of each cache and tune it for better performance (if necessary). See Performance Tuning for more information.

Configurable Caches

System administrators can change the sizes of Confluence's internal caches through the Administration Console and these changes will take effect without the need to first shut down and then restart Confluence. The maximum number of units for any of the defined cache regions can be adjusted individually.

Note that larger cache sizes will require more memory at runtime, so you should review the memory allocation of the Confluence Java process and the physical memory available on your server.

Viewing Cache Statistics and Modifying Cache Sizes

To view the cache statistics:

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Cache Statistics**' in the left-hand panel. There you will find a list of all objects cached within Confluence.
3. Click the '**Advanced**' tab for more detail. Below is an example for one of the most frequently used caches, the 'Content Object' cache.

Name	Percent Used	Effectiveness	Objects / Size	Hit / Miss / Expiry	Adjust Size	Flush
Content Object	80%	73%	4023 / 5000	374550 / 140460 / 55044	*Adjust Size*	*Flush*

About the generated numbers:

Percent Used :	=(Objects)/(Size)
Effectiveness :	=(Hits)/(Hits + Misses)
Objects / Size:	The number of entries in the cache / the number of total possible entries allowed (configurable).
Hit / Miss / Expiry:	The number of reads accessing cache where required content was found / the number of reads accessing cache where required content was not found / the number of objects evicted from the cache.
Adjust Size	Use this option to specify a different maximum cache size. Enter a new cache size and click the '**Adjust Size**' button to set it.
Flush:	Flushes the cache.

For instance, to calculate **Percent Used**:

```
Percent Used = Objects / Size

Percent Used = 4023/5000 = 80%
```

To calculate **Effectiveness**:

```
Effectiveness = (Hits)/(Hits + Misses)

Effectiveness = 374550 / (374550 + 140460) = 73%
```

> ⚠ The clustered versions of Confluence use distributed cache called Tangosol Coherence.

Watching the Cache Contents

To see the specific items in the caches, view the cache statistics at `<baseUrl>/admin/cachecontents.jsp`.

Additional Notes about Configurable Caches

Changes to cache size configurations persist across confluence restarts as they are saved in the `<confluence-home>/config/confluence-coherence-cache-config.xml` file (or `<confluence-home>/config/confluence-coherence-cache-config-clustered.xml` for a clustered instance). In most cases, a Confluence administrator will never need to know about these files. However, if it is necessary to tune cache options other than the maximum cache size, this can be done by manually editing these files. See Cache Performance Tuning for details.

> ⚠️ **Important note about clustered Confluence installations**
> The cache configuration file is stored in a home directory of each cluster node. When a Confluence administrator changes a cache size, all running cluster nodes will automatically update their own configuration files in their respective home directories. However, if a cluster node is not running when an administrator adjusts a cache size, the `/config/confluence-coherence-cache-config-clustered.xml` file in its home directory will not be updated. Since cluster caches are configured by the first node to start, if a node with an outdated cache configuration is the first to start up, the whole cluster would end up using the configuration of that node. However, copying this file from one node to another would resolve this issue.

Performance Tuning

If you need to tune your application when under high usage, you may like to review this document for suggestions.

RELATED TOPICS

Cache Statistics

Cache Performance Tuning for Specific Problems

Cache Performance Tuning

Confluence Cache Schemes

Viewing and Editing License Details

Viewing System Information

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Changing time of Daily Backup

> ⚠️ Atlassian recommends disabling the XML backup both for performance and reliability. XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump. Upgrading is better done without the XML backup. This page can also help with troubleshooting XML Space versions.

By default, Confluence runs its daily backup at 2.00 AM. You can configure Confluence to perform the backup at a time that is best suited to you or your organisational needs.

> ℹ️ **Time is derived from the Confluence server**
> The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:
>
> 1. Go to the Confluence '**Administration Console**'. To do this:
>
>   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
> 2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

Confluence uses Quartz for scheduling periodic jobs. To change the time of your daily backup, you will need to edit the Quartz configuration.

To change the time of your daily backup

1. Open the Quartz configuration file `schedulingSubsystemContext.xml` located under `confluence/WEB-INF/lib/confluence-x.x.x.jar`. Where x.x.x is your Confluence version number.

2. Find the following section of the file:

```
1.<bean id="backupTrigger" class=
"org.springframework.scheduling.quartz.CronTriggerBean">
2.     <property name="jobDetail">
3.         <ref bean="backupJob"/>
4.     </property>
5.     <property name="cronExpression">
6.         <value>0 0 2 * * ?</value>
7.     </property>
8.</bean>
```

3. The string `'0 0 2 * * ?'` sets up a Cron Trigger for the job to run at the zeroth second of the zeroth minute of the 2nd hour, every day of every month, every day of the week.

4. Re-jar the file, either with a zip utility (change the title of .zip back to .jar) or a java command.

5. You can set a new time by editing this string. Note that the date and time format in this configuration file is in this order:

   Second minute hour day

6. Restart Confluence.

For example, to set the new time to twenty past ten PM, change the string to `'0 20 22 * * ?'`.

If you wanted to back up only once a week, for example, at midnight on Sundays, you would change the string to `'0 0 0 ? * SUN'`.

For complete details on the formatting of the cron string, please see http://www.opensymphony.com/quartz/api/org/quartz/CronTrigger.html.

RELATED TOPICS

Alternative Backup Strategy

Site Backup and Restore

Changing time of Daily Backup

Backup FAQ

Configuring Daily Backups

User Submitted Backup & Restore Scripts

Manually Backing Up The Site

Administrators Guide Home Confluence Documentation Home

Confluence Data Directory Configuration

Here is a link listing important Confluence files.

The home directory defines the location of the directory where Confluence will store it's data, including attachments, indexes and backups. Administrators can set this location by defining a value for the file `<MY-INSTALL>/confluence/WEB-INF/classes/confluence-init.properties`. To find what your home directory is currently set to, open this file and check the `confluence.home` property. It is unset on new installations.

Windows Configuration

On Windows, this path:

```
C:\confluence\data
```

will be written like so:

```
confluence.home=C:/confluence/data
```

Note that all backslashes (\) are written as forward slashes (/).

UNIX/Linux/Mac Configuration

On any UNIX-based system, the property is defined using the normal directory syntax:

```
confluence.home=/var/confluence/
```

Symbolic links

If your `confluence.home` directory contains a symbolic link, you must define the absolute path.

> Please note that there can be no symbolic links within the `confluence.home` directory. If disk space is an issue, place the entire `confluence.home` directory on a disk partition where there is enough space.
>
> The absolute path of generated files (such as exports) is compared with the absolute path of the `confluence.home` directory when constructing URLs. When a sub-directory has a different path, the URL will be incorrect, and you may receive "Page not found" errors. These measures are in place to prevent "directory traversal" attacks.

Fixing the Confluence Configuration

The Confluence configuration file: `confluence-cfg.xml` inside the home directory may contain references to the original location of your Confluence home. You will need to edit this file to update these references to also point to the new location. The two properties in this file that need to change are:

- `daily.backup.dir` if you have not configured your backups to be placed elsewhere already
- `hibernate.connection.url` if you are using the embedded HSQL database.

Confluence home directory contents

The Confluence home directory contains data that work in concert with the Confluence database to provide the wiki experience. This document outlines the purpose of the various files and directories in the Confluence home directory.

✅ Tip: Another term for 'Home directory' would be 'data directory'.

Files and directories

confluence.cfg.xml

This file is the most critical file in the Confluence home directory. It contains all of the information necessary for Confluence to start up such as:

- Product license
- Context path
- Database details such as location and connection pool settings
- Paths to important directories

attachments

This directory contains every version of each attachment stored in Confluence. This directory is not used when Confluence is configured to store attachments in the database. Attachments are always stored in the database in clustered instances of Confluence.

Paths within this directory have the following structure:

```
/attachments/PAGE_ID/ATTACHMENT_ID/VERSION
```

An alternative directory may be specified for attachment storage by setting the `attachments.dir` property in `confluence.cfg.xml`

backups

Confluence will place its daily backup archives in this directory, as well as any manually generated backups. Backup files in this directory take the following form:

```
daily-backup-YYYY_MM_DD.zip
```

An alternative directory may be specified for backups by setting the `daily.backup.dir` property in `confluence.cfg.xml`.

bundled-plugins

> ℹ This directory exists for Confluence 2.3 and above

Recent versions of Confluence ship with a set of *bundled* plugins. These are plugins written by the Atlassian and the Confluence community that we think provide useful and broadly applicable functionality in Confluence. The {{bundled-plugins}} directory is where Confluence will unpack its bundled plugins when it starts up. This directory is refreshed on every restart, so removing a plugin from this directory will not uninstall the plugin; it will simply be replaced the next time Confluence starts up.

database

This is where Confluence stores its database when configured to run with the HSQL embedded database and as such contains all Confluence runtime data. Instances configured to run using an external database such as MySQL will not use this directory.

index

This is where Confluence stores its indexes for rapid retrieval of often used data. The Confluence index is used heavily by the application for content searching and recently updated lists and as such is critical for a running Confluence instance. It is important to note however that should the data in this directory be lost or corrupted, it can be restored by running a full reindex from within Confluence. This can take a long time depending on how much data is stored Confluence's database.

An alternative directory may be specified for the index by setting the `lucene.index.dir` property in `confluence.cfg.xml`. As this is the most heavily accessed directory in the Confluence home directory you might want to consider hosting it on the fastest disk available. It would also be useful if the disk holding the Confluence index was not heavily used by any other application to reduce access contention.

plugin-cache

> ℹ This directory exists for Confluence 2.3 and above

As of Confluence 2.3, all Confluence plugins are now stored in the Confluence database. To allow for quicker access to classes contained within the plugin JARs, Confluence will cache these plugins in the `plugin-cache` directory. This directory is updated as plugins are installed and uninstalled from the system and is completely repopulated from the database every time Confluence is restarted. As such, removing plugins from this directory does not uninstall them.

resources

The `resources` directory stores any space logos used in your Confluence instance. For each space with a space logo, there is a directory within `resources` named after the space's key. That directory contains the space's logo.

temp

The `temp` directory is used for various runtime functions such as exporting, importing, file upload and indexing. As the name suggests, and file in this directory is of temporary importance and is only used during runtime. This directory can be safely emptied when Confluence is offline.

An alternative directory may be specified for temporary data by setting the `webwork.multipart.saveDir` property in `confluence.cfg.xml`.

thumbnails

When Confluence generates a thumbnail of an image (for example when the `gallery` macro is used), the resulting thumbnail is stored in this directory for quicker retrieval on subsequent accesses. This directory is essentially a thumbnail cache, and deleting files from this directory simply means the thumbnail will have to be regenerated on the next access.

Confluence 2.2 and older

The following files and directories were used by versions of Confluence older than 2.3

config

The `config` directory is used to store data used by Confluence's bandana data persistence framework. This system is used by Confluence to store the global instance settings and is used by various plugins for their own configuration and data persistence needs. Confluence versions 2.3 and later store these data in the Confluence database and do not use this directory.

The most important file in this directory is the `confluence-global.bandana.xml` file. This file is used to store all of the settings from the `Administration` console in Confluence.

`plugins`

The `plugins` directory is where Confluence stores all installed plugin JARs. It is possible to install and remove plugins by placing and deleting plugin JARs from this directory.

`default-formatting.properties`

This properties file contains various formatting information such as the formats for decimal numbers and dates used in the Confluence user interface. These configuration data where relocated to the Confluence database from Confluence 2.3 and onwards.

Content Index Administration

The Content Indexes power Confluence's search functionality and they are also used for a number of related functions such as building email threads in the mail archive, the Space Activity feature and lists of recently-updated content. The Gliffy Plugin also uses them for some of its functionality.

For reasons of efficiency, content is not immediately added to the index. New and modified Confluence content is first placed in a queue and the queue is processed once every minute (by default).

On this page:

- Viewing the Content Index Summary
- Rebuilding the Content Indexes
- Slow Reindexing
- Viewing the Index Browser
- More Hints and Tips

Viewing the Content Index Summary

To see information about your Confluence instance's content indexing,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.

Screenshot: Index summary

Search Indexes | Queue Contents

Search Index

The search index allows searching of Confluence content.If you are having troubles with search, you may need to rebuild the search index.Please note, rebuilding the search index can severely affect the performance of your instance - it can take hours for some large instances.

BUILT

100% (Elapsed Time: 00:00:12)

Rebuild

Did You Mean Index

You will need to build this index to make "Did You Mean" work. After this has finished, "Did You Mean" will be automatically turned on. **Please note, this feature only provides suggestions for the English language.**

NOT BUILT

0%

Build

Rebuilding the Content Indexes

The content indexes are maintained automatically, but you may need to rebuild one or both of them manually under circumstances such as these:

- Your searching and mail threading are malfunctioning. (Rebuild the Search Index.)
- Your 'Did You Mean' feature is malfunctioning. (Rebuild the Did You Mean Index.)
- After an upgrade. If a content re-index is required after an upgrade, it will be noted in an upgrade subsection of the relevant release notes.

ℹ In new Confluence installations, the 'Did You Mean' feature is not initially activated. To activate it, you first need to build its index by clicking its '**Build**' button on this page.

To rebuild either of the content indexes,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.
3. Click the '**Rebuild**' button in either the 'Search Index' or 'Did You Mean Index' sections on this page, depending on the particular index you want to rebuild.

ℹ
   - If one of these indexes has not yet been built, its button will indicate 'Build' (instead of 'Rebuild').
   - As shown in the image below, only one index can be (re)built at a time.

Screenshot: Content Indexing

Slow Reindexing

Does the reindexing take a long time to complete? The length of time depends on the following factors:

- Number of pages in your Confluence instance.
- Number, type and size of attachments.
- Amount of memory allocated to Confluence.

It may help to increase the heap memory allocation of Confluence by following the instructions in the JIRA documentation.

If you are running an older version of Confluence and find that the index rebuild is not progressing, you may need to shut down Confluence, and restart it with the following *Java system property* set: `bucket.indexing.threads.fixed=1`. This will cause the re-indexing to happen in a single thread and be much more stable (but slower).

Viewing the Index Browser

You can access the Lucene index browser used with Confluence, to see if the required documents have been added to the index.

To view the index browser,

1. Open a browser window.
2. Log in to Confluence.
3. Type the following into the browser address bar: `<Confluence location>/admin/indexbrowser.jsp`.
   For example:

   ```
   http://localhost:8080/admin/indexbrowser.jsp
   ```

4. You will be asked "Which index would you like to browse?"
5. Type in the path of your index directory. By default, this is located at `<Confluence Home directory/index>`.
   For example, `c:\confluence\data\index`.
6. You should see a list of indexed file names.

Screenshot: Index Browser

Which index would you like to browse? `c:\confluence\data\index` [Refresh]

Index Location: c:\confluence\data\index
Number of Documents: 46

1. com.atlassian.confluence.pages.Page-32770 : Thumbnail gallery
2. com.atlassian.confluence.pages.Page-32772 : News
3. com.atlassian.confluence.pages.Page-32773 : Creating pages and linking
4. com.atlassian.confluence.pages.Page-32776 : Index
5. com.atlassian.confluence.pages.Page-32777 : Formatting Content
6. com.atlassian.confluence.pages.Page-32779 : Search Confluence content
7. com.atlassian.confluence.pages.Page-32780 : Using Spaces
8. com.atlassian.confluence.pages.Page-32784 : Email archiving
9. com.atlassian.confluence.pages.Page-32787 : Confluence Overview
10. com.atlassian.confluence.pages.Page-32788 : Jira integration
11. com.atlassian.confluence.pages.Page-32771 : Tasklist
12. com.atlassian.confluence.pages.Page-32774 : Breadcrumb demonstration
13. com.atlassian.confluence.pages.Page-32778 : Source code
14. com.atlassian.confluence.pages.Page-32781 : Example Index
15. com.atlassian.confluence.pages.Page-32785 : RSS Capabilities
16. com.atlassian.confluence.pages.BlogPost-32783 : Octagon blog post
17. com.atlassian.confluence.pages.Attachment-131086 : excelspreadsheet.xls
18. com.atlassian.confluence.pages.Attachment-131089 : office7.jpg
19. com.atlassian.confluence.spaces.Space-65537 :

More Hints and Tips

- If you are still experiencing problems after performing the above rebuild, the next step might be to remove the index and rebuild it from scratch.

  ⚠ The Space Activity feature uses the index to store data. If you remove the index file, the existing activity data will disappear.
- A tip for the development community: If you have the Confluence source, you can look for references to the SmartListManager to find the screens and lists that rely on the content index.

RELATED TOPICS

Creating a Lowercase Page Title Index

Configuring Indexing Language

Rebuild the Content Indices from scratch

Content Index Administration

Working with Macros

[!Administration Guide Attachments directory^adminhome.gif!] 🏠 Confluence Documentation Home

Finding Unused Spaces

Sometimes, you want to know what is *not* being used. It's great to know what's getting most attention, but what about stagnant pages, or even entire spaces that are no longer active?

While viewing space activity and the Global Statistics plugin can provide hints, they still don't always provide enough detail. The simple way is to go directly to the database. We recommend DbVisualizer, and have basic instructions for connecting it to HSQLDB.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
1. SELECT spaces.spacename, MAX(content.lastmoddate)
2. FROM content, spaces
3. WHERE content.spaceid = spaces.spaceid
4. GROUP BY spaces.spacename;
```

It returns a list of spacenames, and the last date and time at which any content was added or changed.

Alternatively, this one simply identifies spaces whose content hasn't changed since a specified date:

```
1. SELECT spaces.spacename
2. FROM content, spaces
3. WHERE content.spaceid = spaces.spaceid
4. GROUP BY spaces.spacename
5. HAVING MAX(content.lastmoddate) < '2006-10-10';
```

The result is a simple list of space names.

It's also possible to present the information in a wiki page, using the SQL plugin, which can be installed via the Plugin Repository. You'll also need to define a database resource in `conf/server.xml` and `confluence/WEB-INF/web.xml`, as described here. Having done so, you can use wiki markup code like the following, replacing `confluenceDS` with the name of your own local datasource:

```
1. h3. Space activity
2. {sql:dataSource=confluenceDS|output=wiki}
3. SELECT spaces.spacename AS Space, MAX(content.lastmoddate) AS LastModified
4. FROM content, spaces
5. WHERE content.spaceid = spaces.spaceid
6. GROUP BY Space;
7. {sql}
```

The result will be something like this:

Space activity:

space	lastmodified
Private Space	2007-10-11 11:34:04.914
Another space	2007-10-11 11:39:39.716
More space	2007-10-11 11:40:11.688

You can try the Chart plugin in combination with the SQL plugin to give more visually attractive results.

Important Directories and Files

The Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.
Important Files and Directories

- `confluence/WEB-INF/classes/confluence-init.properties` : This file tells Confluence where to find the Confluence Home Directory. This file is modified by the administrator when installing Confluence.

- `confluence/WEB-INF/classes/osuser.xml` : This file is modified when connecting Confluence to an external user management system such as an LDAP server or JIRA instance in Confluence 2.0 and earlier. For more information, refer to Understanding User Management in Confluence.

- `confluence/WEB-INF/classes/atlassian-user.xml` : This file is modified when connecting Confluence to an external user management system such as an LDAP server or Crowd. For more information, refer to Understanding User Management in Confluence.

- `confluence/WEB-INF/lib/` : This directory is used when deploying plugins, especially those plugins that cannot automatically be loaded through the Administration Console.

- `confluence/WEB-INF/classes/log4j.properties` : Confluence's logging configuration file. See Working with Confluence Logs.

- `confluence/WEB-INF/classes/ehcache.xml` : This is where you can configure the size of Confluence's internal caches

- `confluence/WEB-INF/classes/styles/site-css.vm` : Confluence's main stylesheet, modify at your own risk

- `conf/server.xml` : SSL configuration.

Memory Settings

The file used to edit JAVA_OPTS memory settings will depend on the method used to install Confluence, as well as the operating system used for your installation.

- Windows Users
  - **Confluence Standalone** — `bin/setenv.bat`
  - **Confluence Installer** — `wrapperwin32.conf`
- Mac/Linux Users
  - **Confluence Standalone** — `bin/setenv.sh`
  - **Confluence Installer** — `wrapperosx.conf`

The Temp Directory

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

Typically, this directory is `/tmp` on Unix systems, or `C:\Temp` on Windows.

To change the location of this directory, you should start the Java Virtual Machine in which confluence is running with the argument:

`-Djava.io.tmpdir=/path/to/your/own/temp/directory.`

The Confluence Home Directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

Tip: Another term for 'Home directory' would be 'data directory'.

Administrators can expect the Confluence Home Directory to grow quite large in a busy site.

The location of this directory is configured by the system administrator during installation (see `confluence-init.properties` above).

Important Files and Directories

- `confluence.cfg.xml` : Confluence's core configuration file; includes the configuration for connecting to its database.

- `default-formatting.properties` : Some auxiliary configuration data concerning default number and date formats.

- `attachments/` : All file attachments in the Confluence site are stored under this directory. This is the only place Confluence keeps attachment files.

- `backups/` : If Confluence is configured to produce daily backups, these are kept in this directory. Administrators should occasionally delete old or unwanted backups from this directory to prevent it from growing too large.

- `config/` : Miscellaneous global and per-space configuration files are kept in this directory.

- `database/` : If Confluence is being run from the embedded HSQL database, the database files will be kept in this directory.

- `index/` : The full-text search index is kept in this directory. Removing or modifying files in this directory may cause search to no longer function. Rebuilding the search index from Confluence's global administration screen will completely regenerate the contents of this directory.

- `plugins/` : Dynamically uploaded plugins are stored in this directory. Administrators can install new plugins by copying them into this directory and triggering a scan from the plugin management page.

- `temp/` : Confluence stores temporary files in this directory, especially during backups and exports. A daily job within Confluence deletes files that are no longer needed.

- `thumbnails/` : Stores temporary files for image thumbnails. The contents of this directory can be safely deleted, as Confluence will regenerate thumbnails as required.

- `velocity/` : Storage for customised page layouts, globally and per-space.

Database

All other data — page contents, links, archived mail and so on — is kept in the database. If you have configured Confluence to use the embedded HSQL database, the database will store its files under `database/` in the Confluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.

> **Tip**
> All of Confluence's persistent data is stored either in the Confluence Home Directory, or the database. If you have backup copies of both of these, taken at the same time, you will be able to restore Confluence from them (see Restoring Data from other Backups).

RELATED TOPICS

Administrators Guide Home [!Administration Guide Attachments directory^dochome.gif!]

Confluence Home Directory

Often in the documentation, you'll see a reference to the 'Confluence Home directory'.

What is the Confluence Home Directory?

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

Tip: Another term for 'Home directory' would be 'data directory'.

You can also read about the contents of the Home directory.

Finding the Confluence Home Directory

The location of the Confluence Home directory is defined when you install Confluence. This location is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.

When Confluence first starts up, it reads the `confluence-init.properties` file to determine where to look for the Home directory.

RELATED TOPICS

Confluence Installation Directory
Important Directories and Files
The Embedded HSQLDB Database

Confluence Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

RELATED TOPICS

Confluence Home Directory
Important Directories and Files

Manually Backing Up The Site

Confluence is configured to make a daily backup of your data and store it as a zipped XML file in the '**backups**' folder under the Confluence Home Directory. A System Administrator can also manually back up the data from the **Administration Console**.

You need to have System Administrator permissions in order to perform this function.

> Consider an alternative backup strategy if your Confluence site is large or you are encountering problems with your automated backup.

To manually back up your site,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Backup & Restore**' in the 'Administration' section of the left-hand panel.
3. Select '**Archive to backups folder**' to store a copy of the backup in the same folder as Confluence's daily backups. (If you do not archive the backup it will be made available for you to download, and then deleted from the server after 24 hours).
4. Select '**Backup attachments**' to include attachments in your backup.
5. Click '**Backup**'.

   Please note that this process will take a few minutes.
6. Once the backup is completed, you will be prompted to download the zipped backup file.

> ✅ If you are running Confluence behind Apache and are facing timeout errors, please consider creating the export directly from Tomcat, instead of going through Apache. This will speed up the process and prevent timeouts.

RELATED TOPICS

Alternative Backup Strategy

Site Backup and Restore

Changing time of Daily Backup

Backup FAQ

Configuring Daily Backups

User Submitted Backup & Restore Scripts

Manually Backing Up The Site

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Configuring Daily Backups

Confluence backs up your data on a daily basis into a zipped XML file. By default, the backup is performed at 2.00 a.m. and the backup files are stored in the `backups` folder under the Confluence Home Directory.

> 🛈 **Time is derived from the Confluence server**
> The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:
>
> 1. Go to the Confluence '**Administration Console**'. To do this:
>
>   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
> 2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'. You can change the time of the daily backup.

> ✅ Consider an alternative backup strategy if your Confluence site is large or you are encountering problems with your automated backup.

The default naming convention for the backup files is 'daily-backup-yyyy_MM_dd'. Confluence can write backups to both local and mapped network drives.

From the **Administration Console**, you can:

- Enable or disable backups.
- Include or exclude attachments in backups.
- Configure a different path to store backup files.
- Change the naming format used for the files.

🛈 You need to have System Administrator permissions in order to perform this function.

To configure your daily backups,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Daily Backup Admin**' in the '**Configuration**' section.
3. Click the '**Edit**' button on the '**Daily Backup Administration**' screen.
4. Now you can do the following:
    - To disable backups — Select '**Disable**'.
    - To specify an alternate path to store backup files — Select '**Custom**' and then enter the path. The directory must be on either a local drive or a mounted network drive.

      ⓘ Please ensure the mapped drive is on a physical server, not a Virtual Machine image.
    - To exclude attachments from backups — Select '**Off**' beside '**Backup Attachments**'. By default, this is 'On'.
    - To use a different naming prefix format — Enter the new format in the '**Backup File Prefix**' input field.
    - To use a different date format — Enter the date format in the '**Backup File Date Pattern**' input field using the syntax described in this document from Sun.
5. '**Save**' your changes.

Below is an example of daily backup being *disabled*.

Backup Path	○ **Default** (/opt/j2ee/domains/atlassian.com/confluence/webapps/atlassian-confluence/data/backups) ○ **Custom** ◉ **Disable**
Backup Attachments:	◉ **On** ○ **Off**
Backup File Prefix:	daily-backup-
Backup File Date Pattern:	yyyy_MM_dd

Save

RELATED TOPICS

Alternative Backup Strategy

Site Backup and Restore

Changing time of Daily Backup

Backup FAQ

Configuring Daily Backups

User Submitted Backup & Restore Scripts

Manually Backing Up The Site

Administrators Guide Home   Confluence Documentation Home

User Submitted Backup & Restore Scripts

These scripts are user-submitted and should be used with caution as they are not covered by Atlassian technical support. If you have questions on how to use or modify these scripts, please post them to the Confluence user forum. Feel free to submit new scripts or post updates by logging in and adding them to the page as a comment.

Delete Old Backups - Wscript Script On Windows

This script examines backup filename and deletes them if necessary, it may need to be edited.

```
01.'If you want 3 day old files to be deleted then insert 3 next to Date - "your number here"
02.'This script will search out and delete files with this string in them ".2005-12-04-" This of
course depends on the number you enter.
03.'You can always do a wscript.echo strYesterday or strFileName to see what the script thinks you
are searching for.
04.
05.dtmYesterday = Date - 3
06.
07.strYear = Year(dtmYesterday)
08.
09.strMonth = Month(dtmYesterday)
10.If Len(strMonth) = 1 Then
11.          strMonth = "0" & strMonth
12.End If
13.
14.strDay = Day(dtmYesterday)
15.If Len(strDay) = 1 Then
16.          strDay = "0" & strDay
17.End If
18.
19.strYesterday = strYear & "-" & strMonth & "-" & strDay
20.
21.strFileName = "C:\test*." & strYesterday &"-*"
22.
23.Set objFSO = CreateObject("Scripting.FileSystemObject")
24.objFSO.DeleteFile(strFileName)
```

Delete Old Backups - Basic Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following:

```
1.ls -t <path to your backup dir>/* | tail -n +6 | xargs -i rm {}
```

Or, using the older form of the `tail` command if your system does not support the standard form:

```
1.ls -t <path to your backup dir>/* | tail +6 | xargs -i rm {}
```

Delete Old Backups - Advanced Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following. Set the BACKUP_DIR and DAYS_TO_RETAIN variables to appropriate values for your site. Between runs, more files than DAYS_TO_RETAIN builds up.

```
01.#!/bin/sh
02.
03.# Script to remove the older Confluence backup files.
04.# Currently we retain at least the last two weeks worth
05.# of backup files in order to restore if needed.
06.
07.BACKUP_DIR="/data/web/confluence/backups"
08.DAYS_TO_RETAIN=14
09.
10.find $BACKUP_DIR -maxdepth 1 -type f -ctime +$DAYS_TO_RETAIN -delete
```

Manual Database & Home Backup - Bash Script For Linux

This backs up a mySQL database and the Confluence home directory.

```
1.#!/bin/bash
2.CNFL=/var/confluence
3.CNFL_BACKUP=/backup/cnflBackup/`date +%Y%m%d-%H%M%S`
4.
5.rm -rf $CNFL/temp/*
6.mkdir $CNFL_BACKUP
7.mysqldump -uroot -p<password> confluence|gzip > $CNFL_BACKUP/confluence.mysql.data.gz
8.tar -cjvf $CNFL_BACKUP/data.bzip $CNFL > $CNFL_BACKUP/homedir.status
```

Backup by Date - Postgres

```
1.export d=`date +%u`
2.mkdir -p /home/backup/postgres/$d
3.
4.sudo -u postgres pg_dumpall | bzip2 > /home/backup/postgres/$d/sql.bz2
```

Related Topics

- Site Backup and Restore
- Backup FAQ

Migrating Confluence Between Servers

> Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their `{recently-updated}` macro is not being updated correctly. Errors in the `atlassian-confluence.log` file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a rebuild of your content indices after performing a migration.

How to Create a Test or Development Instance

> Development licenses are available for any Commercial or Academic license. Create one or contact us for help.

Adminstrators may need to move a Confluence instance from one server to another for upgrades or downtime. This page tells you how to copy a Confluence instance from one server to another. For example, you may want to transfer your current production snapshot to a test server as permitted in the licence agreement.

> **Avoid upgrades while transferring**
> If you are planning to switch databases, application servers or Confluence versions, perform the transfer and test that it is successful **separately** to any other changes.

Transferring Confluence To Another Server Using The Same Operating System

If the operating systems on both servers are the same, then the home and install folders can be copied straight into an identical external database and user management setup.

1. On the original server, create zips of the Confleunce install and home directories. Copy the zips to the new server.

2. On the new server, unzip the install and home directories. Windows users should avoid unzipping with the Windows built-in extractor, instead use Winzip or the free 7Zip.

   > If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.

3. Modify the location of your war file if need be. If using Tomcat, this is likely in /Conf/Catalina/localhost. You'll want to make sure the docbase attribute is pointing to the right location.
4. This next step is dependent on your database:

   a. For users of the internal database, the content is stored inside the home directory. You should switch to an external database after the transfer is successful.
      i. For databases stored on another server: change the user account or datasource permissions so that the new server has the same network access permissions as the original. Then confirm from the new server that the hostname can be resolved and is listening for database connections on the expected port.
      ii. For external databases stored locally: on the original server, create a manual database backup using a native db dump backup tool. Copy the database backup to the new server.
   b. On the new server, install or upgrade the database version to match the original server.
   c. Import the database backup.
   d. Add a database user account with the same username and password as the original.
   e. Provide the user with the full access to the imported database.
   f. Use a database administration tool to confirm that the user can login from the localhost.
   g. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
   h. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test instance, set up a test of your JIRA instance or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
   i. If appropriate, make sure no emails are sent out from the test system.
   j. Start Confluence.
   k. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
   l. If you configured Confluence as a Windows service, repeat those instructions.
   m. Add your development license key.

Transferring Confluence To Another Server Using a Different Operating System

Using database tools (preferred option)

If you are using the Alternate backup strategy, follow these steps:

1. Download the proper distribution (**the same one you have from your original instance**) from the Download Archive.
2. Copy your Confluence *home* (not install) directory from your original server (even if it was a different OS).
3. If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.
4. For external databases stored locally, on the original server, create a manual database backup using a native db dump backup tool.
5. Copy the database backup to the new server.
6. On the new server, install or upgrade the database version to match the original server.
7. Import the database backup.
8. Add a database user account with the same username and password as the original.
9. Provide the user with the full access to the imported database.
10. Use a database administration tool to confirm that the user can login from the localhost.
11. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
12. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original.
13. Copy server.xml, atlassian-user.xml, osuser.xml, any patches, and any other customized files velocity or properties files. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test instance, set up a test of your JIRA instance or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
14. If appropriate, make sure no emails are sent out from the test system.
15. Start Confluence.
16. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
17. If you configured Confluence as a Windows service, repeat those instructions.
18. Add your development license key.

For XML backups (only for small to medium sized installations)

If you're not yet using the Alternate backup strategy, you can do this with your regular XML backup. Create a backup and import into the new server.

1. Create a backup from Confluence:

    a. Go to the Confluence '**Administration Console**'. To do this:

        • Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
    b. Select '**Backup & Restore**'.
    c. Check the '**Backup Attachments**' option and select '**Backup**'.

2. Identify the current version of Confluence your are using, displayed at the bottom of each Confluence page.
3. Download the same version as you are currently using to the new server, which may be the current Confluence release, or an older version.
4. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
5. Using the same version, follow the Upgrading Confluence guide.
6. Add your development license key.
7. Restore your XML Backup From <<Administration > Backup and Restore>>.
8. If appropriate, make sure no emails are sent out from the test system.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

    ```
    1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.smtp.mail.accounts';
    ```

2. Disable space-level mail archiving by running the following database query:

    ```
    1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.space.mailaccounts';
    ```

Change the 'SELECT *' to a 'DELETE' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test instance without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, JIRA macro, etc.). I They write content, or create unwanted load on external systems, they should be disabled promptly after starting the test instance.

> **Blog post on Moving Confluence from Windows to Linux**
> Ricky Sheaves (calebscreek) has written an interesting blog post on Moving Confluence from Windows to (Ubuntu) Linux.

Merging instances

If you wish to merge two instances, you can consider using the remote import plugin. This plugin is currently unsupported. The supported method would be to export a space and then import spaces one by one. The two instances must be the same version.

Rebuilding the Ancestor Table

In Confluence, the ancestor table controls the breadcrumb navigation at the top of each Confluence page. Occassionally, the ancestor table will become out of sync. When this happens, you can rebuild the table to restore everything to normal.

Simply access this URL:

```
http://yoursite/admin/permissions/pagepermsadmin.action
```

Screenshot: Page Level Permissions

Dashboard > Administration > Page Level Permissions

Administration
Page Level Permissions

Configuration

☐ General Configuration
☐ Daily Backup Admin
☐ Manage Referrers
☐ Plugins

Rebuild Ancestor Table
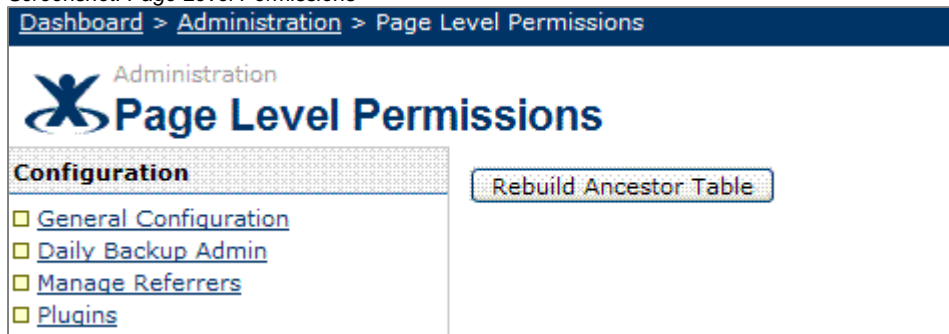
RELATED TOPICS

Administrators Guide Home   Confluence Documentation Home

Restoring a Site

⛔ CAUTION: Restoring a backup of an entire confluence site (consisting of multiple spaces) will:

- Wipe out all Confluence content in the database. Ensure that your database is backed up.
- Log you out after the restore process. Make sure you know your login details contained in the data being restored.

⚠ Atlassian suggests establishing a backup strategy using a native database tool for a production instance of Confluence.

Confluence supports **backward compatibility** for site backups. (But **not** for space backups). You can only successfully restore backups of a site from an older version of Confluence to a newer version of Confluence. For example, if you create a site backup in Confluence 2.4.3, it cannot be restored into a Confluence 2.2.2 instance. It can however, be restored into 2.4.5 or 2.5.x, because 2.4.5 and 2.5.x are newer versions of Confluence.

There are two ways to restore a site from a backup file:

1. Restore a site from the Confluence Setup Wizard: This restores the data into a new instance of Confluence.
2. Restore a site from the Administration Console: This restores data into the current instance of Confluence.

If your daily backup zips cannot be restored for whatever reason, but you have backups of both your database and your Confluence home directory, then it is still possible to restore from these backups.

⚠ **Selective space restore not possible**
You cannot select a single space to restore from the entire site backup when the backup contains more than one space.

RELATED TOPICS

📄 Restoring from Backup During Setup

Restoring Data from the Administration Console

Restoring a Space

Restoring a Site

Manually Backing Up The Site

Confluence Docs 3.1

Administrators Guide Home   Confluence Documentation Home

Restoring a Space

This page tells you how to import the contents of a Confluence space into another Confluence site, via an XML backup file.

You can export the content of a space, including pages, comments and attachments. The process involves converting the data in the space into XML format. The end product is a zip file that contains XML file(s) and optionally, all the attachments in the space. To transfer this data to another Confluence site, you simply restore this zip file as described below.

Confluence will only allow you to restore a space if there is not already a space by that name on the site. If you already have a space with the identical name, you will need to delete or rename the existing space before restoring the new one.

> 🚫 **Cannot restore to a different major Confluence release**
> Confluence only supports **forward compatibility** and **backward compatibility** for individual space import and export when executed within the **same major version** of Confluence instances.
>
> **Restoration Data Must Share the Same Major Version Number**
> This means that a space export created in a newer major version of Confluence cannot be imported into an older major version of Confluence. For example, if you create a space export in Confluence 2.4.5, it cannot be imported into a Confluence 2.2.2 instance. It can be however imported into 2.4.6. (because 2.2.2 and 2.4.5 are two different 'major' versions). Similarly, a space export created in 2.2.2 can not be imported into 2.4.5. However, it can be restored in 2.2.10 (since 2.2.2 and 2.2.10 belong to the same major version release).
>
> If such an operation is carried out, an error message similar to the one below will be displayed and the import action will be stopped.
>
> *Screenshot: Major Version Clash on Space Restore*
>
> The following error(s) occurred:
> > ● Restore denied. You can only restore space backups exported from the same major version (e.g. 2.2.x or 2.3.x).
>
> You'll need to set up a test server of the same version and import the space, then upgrade your test installation so it's the right major version so that you can perform the export and import successfully. Otherwise, you can try to Change the version of the space export, but please try this on a test instance as well.

ℹ️ You need to have System Administrator permissions in order to perform this function.

To restore a space,

> 1. Go to the Confluence '**Administration Console**'. To do this:
>
>   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
> 2. Select '**Backup and Restore**' in the 'Administration' section of the left-hand panel.
>
> You can restore data in one of two ways:
>
> 1. **Upload a zipped backup to Confluence:**
>   - Browse for the backup file.
>   - Uncheck '**Build Index**' if you want to create the index at a later stage.
>   - Click '**Upload and Restore**'.
>
> 2. **Restore a backup from the file system:**
>   - Select the backup file from the form field displayed. If you do not see your backup file, make you sure that it has been copied into the `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore` directory.
>   - Uncheck '**Build Index**' if you want to create the index at a later stage.
>   - Click '**Restore**'.

Workaround for restoring Spaces between Major Releases

As described in the warning above, you cannot restore a space directly between two major versions. i.e. an individual space backup from 2.2.x cannot be imported into a 2.3.x.
Whilst there is presently no functionality in Confluence that allows you to do so, there is however a manual workaround to this problem.

⚠ **Always back up your data** before attempting this procedure.

Please follow the procedure below:

1. Create a manual site backup
2. Configure a separate installation of the Confluence version you want the space export for i.e. the version into which you want to export the space.
3. Use one of the two restoration methods listed here to import the site backup.
4. After restoring the site backup, create the XML backup for the space you want.
5. Import this backup into your main Confluence instance and you now have that space.

Alternative workaround

Another way is to change the version of a space backup.

RELATED TOPICS

📄 Restoring from Backup During Setup

📄 Restoring Data from the Administration Console

📄 Restoring a Space

📄 Restoring a Site

📄 Manually Backing Up The Site

🌐 Confluence Docs 3.1

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Changing the version of a space backup

Confluence prevents the import of space backups which aren't from the same major version. The reason for this is that any schema change between the export and imported version of Confluence will cause the import to fail, leaving you with an incomplete import. Even worse, the failure can be database-dependent, so it may work fine on one particular database but your backup will fail to import later.

> ⛔ Do not import a modified space backup on a production server. Import the modified space backup on a test server, then export from the test server to create a pristine space backup for the new version.

To change the version of a space backup, do the following:

* extract the space backup ZIP file
* edit `exportDescriptor.properties` in a text editor
* change the buildNumber to the buildNumber of the Confluence version you wish to import into
* zip up the modified contents of the backup into a ZIP file again.

This will allow you to import a backup into a test instance of Confluence. After checking the imported space for errors, export it cleanly from the test server and import the fresh backup into your production server.

If your import fails on the test server due to Hibernate errors, this indicates a schema incompatibility and cannot be worked around. You will need to restore your entire site on an old version of Confluence, and export the space from there. See the last section of Restoring a Space for details.

Restoring a Test Instance from Production

Many Confluence administrators will have a production instance running the "live" version of Confluence, as well as a test instance for testing upgrades and so on. In this situation, it's quite common that the two instances are running different versions of Confluence. This document describes how to copy the data from a production instance to a test instance, where the production version may be different to the test version.

Before proceeding with this guide, ensure you have read and understood the normal procedure for upgrading Confluence.

Upgrading a test Confluence instance with production data

Essentially, we are copying both the production home directory and database to the test instance. We then update the database details on the test instance to point to the test database, leaving all other instance metadata (most importantly the Confluence *build number*) the same as production.

1. Shut down your test instance.
2. Restore the production database to the test database server.
3. Create a backup of the `confluence.cfg.xml` file found in the home directory of the test instance.
4. Copy the production confluence-home directory to the test application server.
5. Open the `confluence.cfg.xml` which has been copied in a text editor. Change the database settings to match the test database server. **Ensure you do not point to your production database.** (You can compare with the backup you made in Step 3 if you need to get the database settings. Don't just copy this file – you need the build number unchanged from production to indicate the database is from an older version of Confluence.)

Before starting your test instance, you need to do the following steps to ensure no contact with production systems.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

   ```
   1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.smtp.mail.accounts';
   ```

2. Disable space-level mail archiving by running the following database query:

   ```
   1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.space.mailaccounts';
   ```

Change the 'SELECT *' to a 'DELETE' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test instance without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, etc.). These should be disabled promptly after starting the test instance.

You can create a developer license for this server and update the License Details after starting up.

See also

Upgrading Confluence
Migrating Confluence Between Servers
Restoring to a Test Instance of Confluence from Production

Restoring Data from other Backups

Typically, Confluence data is restored from the Administration Console or from the Confluence Setup Wizard.

If you are experiencing problems restoring from an zipped XML backup file, it is still possible to restore provided you have:

1. A backup of your home directory.
2. A backup of your database (if you're using an external database).

Instructions for this method of restoring differ depending on whether you are using the embedded database or an external database (like Oracle, MS SQL Server, MySQL or Postgres).

Embedded Database

If you are running against the embedded database, the database is located inside the `database` folder of your Confluence Home Directory. Hence, all you need to do is:

1. Retrieve the most recent backup of your home directory.
2. Unpack the Confluence distribution and point the `confluence-init.properties` file to this directory.

External Database

If you're using an external database, you need to do the following.

1. Prepare backups of your home directory and database (preferably backups that are dated the same). That is, make sure the home directory is accessible on the filesystem and the database available to be connected to.
2. If this database happens to have a different name, or is on a different server, you need to modify the jdbc url in the `confluence.cfg.xml` file inside the Confluence Home Directory. The value of this property is specified as `hibernate.connection.url`.
3. Unpack the Confluence distribution and point the `confluence-init.properties` file to the home directory.

RELATED TOPICS

Important Directories and Files

Migrating to a Different Database

Restoring Data from the Administration Console

Use this option if you want to restore data into your current instance of Confluence. If you want to restore data into a new instance, follow the instructions here.

ⓘ You need to have System Administrator permissions in order to perform this function.

⊖ CAUTION: Restoring a backup of an entire Confluence site (consisting of multiple spaces) will do the following:

- Wipe out all Confluence content in the database. Ensure that your database is backed up.
- Log you out after the restore process. Make sure you know your login details contained in the data being restored.

To restore data from backup,

- Go to the Confluence '**Administration Console**'. To do this:

  - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
- Select '**Backup and Restore**' in the 'Administration' section of the left-hand panel.

You can restore data in one of two ways:

1. **Upload a zipped backup to Confluence:**
   - Browse for the backup file.
   - Uncheck '**Build Index**' if you want to create the index at a later stage.
   - Click '**Upload and Restore**'.
2. **Restore a backup from the file system:**
   - Select the backup file from the form field displayed. If you do not see your backup file, make sure that it has been copied into the
     `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore` directory.
   - Uncheck '**Build Index**' if you want to create the index at a later stage.
   - Click '**Restore**'.

RELATED TOPICS

📄 Restoring from Backup During Setup

📄 Restoring Data from the Administration Console

📄 Restoring a Space

📄 Restoring a Site

📄 Manually Backing Up The Site

🌐 Confluence Docs 3.1

Retrieve file attachments from a backup

File attachments on pages can be retrieved from a backup without needing to import the the backup into Confluence. This is useful for recovering attachments that have been deleted by users.

Both daily and manual backups allow this, as long as the 'Include attachments' property was set. Users wanting to restore pages, spaces or sites should check out the Administrators Guide instead.

Before following the instructions for recovering attachments, please review how backups store file and page information.

How Backups Store File and Page Information

The backup zip file contains entities.xml, an XML file containing the Confluence content, and a directory for storing attachments.

Backup Zip File Structure

Page attachments are stored under the attachments directory by page and attachment id. Here is an example listing:

```
Listing for test-2006033012_00_00.zip
\attachments\98\10001
\attachments\98\10002
\attachments\99\10001
entities.xml
```

Inside the attachment directory, each numbered directory inside is one page, and the numbered file inside is one attachment. The directory number is the page id, and the file number is the attachment id. For example, the file \attachments\98\10001 is an attachment with page id 98 and attachment id 10001. You can read entities.xml to link those numbers to the original filename. Entities.xml also links each page id to the page title.

Entities.xml Attachment Object

Inside the entities.xml is an Attachment object written in XML. In this example, the page id is 98, the attachment id is 10001 and the filename is myimportantfile.doc. The rest of the XML can be ignored:

```
1.<object class="Attachment" package="com.atlassian.confluence.pages">
2.<id name="id">98</id>
3.<property name="fileName"><![CDATA[myimportantfile.doc]]></property>
4....
5.<property name="content" class="Page" package="com.atlassian.confluence.pages"><id name="id"
>10001</id>
6.</property>
7....
8.</object>
```

Entities.xml Page Object

This XML describes a page. In this example, the page id is 98 and the title is Editing Your Files. The rest of the XML can be ignored:

```
1.<object class="Page" package="com.atlassian.confluence.pages">
2.<id name="id">98</id>
3.<property name="title"><![CDATA[Editing Your Files]]></property>
4....
5.</object>
```

Instructions for Recovering Attachments

Each file must be individually renamed and re-uploaded back into Confluence by following the instructions below. Choose one of the three methods:

Choice A - Recover Attachments By Filename

Best if you know each filename you need to restore, especially if you want just a few files:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the filename and find the attachment object with that filename. Locate its page and attachment id.
3. Using the page and attachment id from entities.xml, go to the attachments directory and open that directory with that page id. Locate the file with the attachment id.
4. Rename the file to the original filename and test it.
5. Repeat for each file.
6. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice B - Restore Files By Page

Best if you only want to restore attachments for certain pages:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the page title and find the page object with that title. Locate its page id.
3. Go to the attachments directory and open that directory with that page id. Each of the files in the directory is an attachment that must be renamed.
4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id and filename.
5. Rename the file with that attachment id to the original filename and test it.
6. Repeat for each page.
7. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice C - Restore All Files

Best if you have a small backup but want to restore many or all the attachments inside:

> 🚫 Following process is applicable to **space** export only. Site xml backups do not require page id to be updated manually due to the nature of persistent page_id's.

1. Unzip the backup directory and open entities.xml.
2. Go to the attachments directory and open any directory. The directory name is a page id. Each of the files in the directory is an attachment that must be renamed.
3. Search entities.xml for attachment objects with that page id. When one is found, locate the attachment id and filename.
4. Rename the file with that attachment id to the original filename and test it.
5. Find the next attachment id and rename it. Repeat for each file in the directory.
6. Once all files in the current directory are renamed to their original filenames, search entities.xml for the page id, eg directory name. Find the page object with that page id and locate its page title.
7. Rename the directory to the page title and move on to the next directory. Repeat for each un-renamed directory in the attachments directory.
8. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

> ⚠️ To obtain detailed information about lost attachments, **location**, **name** and **type** of the attachments, you may use the [findattachments script]

Troubleshooting failed XML site backups

> ⚠️ XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

Resolve Errors With Creating An XML Backup

The errors may be caused by a slightly corrupt database. If you're seeing errors such as 'Couldn't backup database data' in your logs, this guide will help you correct the error on your own. We strongly recommend that you backup your database and your Confluence home directory beforehand, so that you can restore your site from those if required. If you are unfamiliar with SQL, we suggest you contact your database administrator for assistance.

Preferable solution

The alternative backup strategy is a very reliable and often more efficient way to do backups. If you are running into problems with XML backups - whether memory related or because of problems like the one described here - consider using a native backup tool as an alternate solution.

To Identify And Correct The Problem

To work out where the data corruption or problems are, increase the status information reported during backup, then edit the invalid database entry:

1. Stop Confluence.
2. If you have an external database, use a database administration tool to create a manual database backup.
3. Backup your Confluence home directory. You will be able to restore your whole site using this and the database backup.
4. Open the `my_confluence_install/confluence/WEB-INF/classes/log4j.properties` and add this to the bottom and save:

   ```
   1.log4j.logger.com.atlassian.confluence.importexport.impl.XMLDatabinder=DEBUG, confluencelog
   2.log4j.additivity.com.atlassian.confluence.importexport.impl.XMLDatabinder=false
   ```

5. Find your atlassian-confluence.log. Move or delete all existing Confluence logs to make it easier to find the relevant logging output.
6. Restart Confluence and login.
7. Begin a backup so that the error reoccurs.
8. You must now check your log files to find out what object could not be converted into XML format. Open `my_confluence_install/logs/catalina.out`. Scroll to the bottom of the file.
9. Do a search for 'ObjectNotFoundException'. You should see an error similar to this:

```
01 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder]
   Writing object: com.atlassian.confluence.core.ContentPermission with ID: 5 to XML.
02 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder]
   Writing property: type
03 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder]
   Writing property: group
04 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder]
   Writing property: expiry
05 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder]
   Writing property: content
06 [DOCPRIV2:ERROR] LazyInitializer - Exception initializing proxy
   <net.sf.hibernate.ObjectNotFoundException: No row with the given identifier exists: 2535,
07 of class:
   com.atlassian.confluence.core.ContentEntityObject>net.sf.hibernate.ObjectNotFoundException:
08 No row with the given identifier exists: 2535, of class:
   com.atlassian.confluence.core.ContentEntityObject
09 at net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFoundException.java:24)
10 at net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java:1946)
11 at net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitializer.java:53)
12 at
   net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(LazyInitializer.java:60)
13 at net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyInitializer.java:164)
14 at net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyInitializer.java:108)
15 at
   com.atlassian.confluence.core.ContentEntityObject$$EnhancerByCGLIB$$cc2f5557.hashCode(<generate
   at java.util.HashMap.hash(HashMap.java:261)
17 at java.util.HashMap.containsKey(HashMap.java:339)
18 at
   com.atlassian.confluence.importexport.impl.XMLDatabinder.toGenericXML(XMLDatabinder.java:155)
```

10. Open a DBA tool such as DbVisualizer and connect to your database instance. Scan the table names in the schema. You will have to modify a row in one of these tables.
11. To work out which table, open `catalina.out`, check the first line of the exception. This says there was an error writing the `ContentPermission` object with id 5 into XML. This translates as *the row with primary key 5 in the CONTENTLOCK table* needs fixing. To work out what table an object maps to in the database, here's a rough guide:
    * Pages, blogposts, comments --> CONTENT table
    * attachments --> ATTACHMENTS table
    * More information can be found in the schema documentation
12. Now you must find the primary key of the incorrect row in this table. In this case, you can check the first line and see that the row has a primary key of 5.
13. Each property is written to a column, so the last property that was being written has the incorrect value. The row being written to when the exception was thrown was CONTENT (line 5) with a value of 2535 (line 6). Now you know the column and value. This value 2535 is the id of an entry that no longer exists.
14. Using a database administrative tool, login ot the Confluence database. Locate the row in the relevant table and correct the entry. Check other rows in the table for the default column value, which may be null, 0 or blank. Overwrite the invalid row value with the default.
15. Restart Confluence.
16. Attempt the backup again. If the backup fails and you are stuck, please lodge a support request with your latest logs.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.BucketPropertySetItem@a70067d3
SQL []; Violation of PRIMARY KEY constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate
key in object 'OS_PROPERTYENTRY'.; nested exception is java.sql.SQLException: Violation of PRIMARY
KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM OS_PROPERTYENTRY GROUP BY
ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING COUNT(*)>1
```

To Help Prevent This Issue From Reoccuring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

RELATED TOPICS

Enabling detailed SQL logging

🏠Administrators Guide Home 🏠Confluence Documentation Home

Troubleshooting XML backups that fail on restore

⚠️ XML site backups are only necessary for migrating to a new database. Upgrading Confluence, Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

ℹ️ If migrating from HSQLDB to MySQL, you might have a better experience using the MySQL Migration Toolkit.

Seeing an error when creating or importing a site or space backup?

Problem	Solution
Exception while creating backup	Follow Troubleshooting failed XML site backups instead
Exception while importing backup	Follow instructions below

Resolve Errors When Attempting To Restore An XML Backup

The errors may be caused by a slightly corrupt database. You will need to find the XML backup file entry that is violating the DB rules, modify the entry and recreate the XML backup:

1. On the instance being restored, follow the instructions to disable batched updates (for simpler debugging), log SQL queries and log SQL queries **with parameters** at Enabling detailed SQL logging.
2. Once all three changes have been made, restart Confluence.
3. Attempt another restore.
4. Once the restore fails, check your log files to find out what object could not be converted into XML format. For Confluence Standalone users, check your Confluence install directory under the `/logs/` and check both `atlassian-confluence.log` and `catalina.out` file. The correct file will contain SQL debug output.
5. Scroll to the bottom of the file and identify the last error relating to a violation of the database constraint. For example:

```
2006-07-13 09:32:33,372 ERROR [confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.exception.ConstraintViolationException:
  could not insert: [com.atlassian.confluence.pages.Attachment#38]
net.sf.hibernate.exception.ConstraintViolationException: could not insert:
[com.atlassian.confluence.pages.Attachment#38]
...
Caused by: java.sql.SQLException: ORA-01400: cannot insert NULL into
("CONFUSER"."ATTACHMENTS"."TITLE")
at oracle.jdbc.driver.DatabaseError.throwSqlException(DatabaseError.java:112)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:331)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:288)
```

This example indicates a row in your attachment table with ID = 38 that has a null title.

6. Go to the server that the backup was created on. You must have a copy of the database from which the backup was created. If you do not have this, use a DBA tool to restore a manual backup of the database.
7. Open a DBA tool and connect to the original database instance and scan the table names in the schema. You will have to modify a row in one of these tables.
8. To work out which table, open `catalina.out`, check the first line of the exception. To work out what table an object maps to in the database, here's a rough guide:
   - Pages, blogposts, comments --> CONTENT table.
   - attachments --> ATTACHMENTS table.
9. To correct the example error, go to the attachment table and find that attachment object with id 38. This will have a a null title. Give a title using the other attachments titles as a guide. You may have a different error and should modify the database accordingly.
10. Once the entry has been corrected, create the XML backup again.
11. Import the backup into the new version.
12. If the import succeeds, revert the changes made in your SQL logging to re-enable disable batched updates and turn off log SQL queries and log SQL queries with parameters.

13. Restart Confluence.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.BucketPropertySetItem@a70067d3
SQL []; Violation of PRIMARY KEY constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate
key in object 'OS_PROPERTYENTRY'.; nested exception is java.sql.SQLException: Violation of PRIMARY
KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

This indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM OS_PROPERTYENTRY GROUP BY
ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING COUNT(*)>1
```

Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems

If you're receiving a message like:

```
ERROR [Importing data task] [confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.PropertyValueException: not-null property references a null or transient value:
com.atlassian.user.impl.hibernate.DefaultHibernateUser.name
```

This means there's an unexpected null value in a table. In the above example, the error is in the name column in the USERS table. We've also seen them in the ATTACHMENTS table.

Remove the row with the null value, redo the xml export, and reimport.

To Help Prevent this Issue from Reccuring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

> ⚠ The problem with different settings for case sensitivity varies between databases. The case sensitivity of the database is usually set through the collation that it uses. Please vote on the existing issue

RELATED TOPICS

Troubleshooting failed XML site backups
Administrators Guide

Migrating from HSQLDB to MySQL

> ⚠ If you've gone through Migrate to Another Database and cannot migrate because of a failed xml backup, this page might help.

Disclaimer

MySQL Migration Toolkit is released by the makers of MySQL and as such, problems with the software should be directed to them. Atlassian Support does not offer support for the Migration Toolkit, nor do we provide support for this migration path. These instructions are offered for strictly informational purposes, and your mileage may vary.

> ⛔ **Backup Reminder**
> Please backup your database and your home folder before attempting this.

Resources needed:

- Empty MySQL DB with appropriate credentials to allow creation, deletion, and insertion of tables and rows.
- A Windows machine that can both communicate to the Confluence server and the destination DB.
- MySQL Migration Toolkit
- HSQL Database Engine

Preparation for migrating to MySQL from HSQLDB

1. Shutdown Confluence
2. Make a copy of the confluence home folder for backup purposes
3. Install the Migration Toolkit
4. Unzip the hsqldb package.
5. Copy the hsqldb.jar from hsqldb/lib into C:\Program Files\MySQL\MySQL Tools for 5.0\java\lib
6. Start the MySQL Migration Toolkit

Running the Migration Toolkit

You should be presented with the following screen.
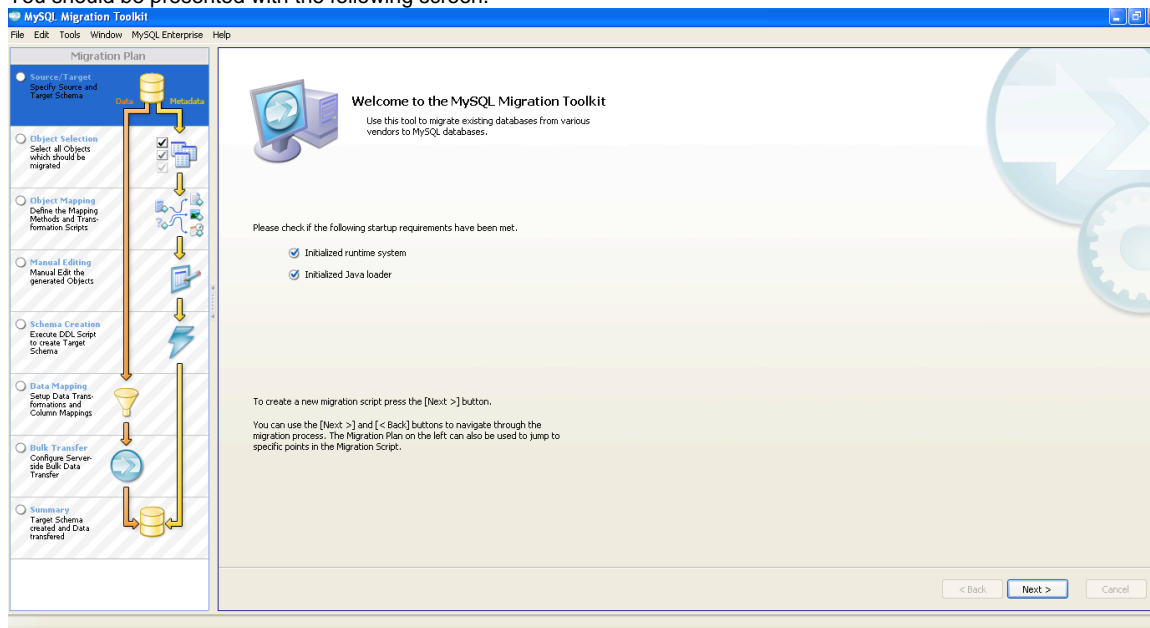
Choose Direct Migration

Source Database

Database System:	Generic JDBC
Class Name:	jdbc:hsqldb: file:PATHTODATABASEFOLDER\confluencedb
Username:	sa
Password:	*No password. Leave this field blank*

Destination Database

🚫 Please make sure that the computer that is running MySQL Toolkit is able to access the MySQL server and that the user listed has the ability to create, drop, insert, and update tables.

Connecting to Servers

Connecting to Servers
Establishing database connections.

Connection Progress

Tasks to execute
The following tasks will now be executed. Please monitor the execution progress. Press [Advanced >>] to see the log.

☑ Connecting to source database system

☑ Retrieve schema information from source database system

☑ Test connection to target database system

Execution completed successfully.

You should see the toolkit trying to connect. If you have problems, please click on the advanced options and sql will show you debugging information. Click Advanced to see the log. If you see "Java Heap Space: Out of Memory", you can start the *MySQL Migration Toolkit* with a -Xmx flag to allocate more memory to the JVM.

After this screen you should come to reverse engineering. Click next.

Source Schemata Selection

You should see 2 databases, **INFORMATION_SCHEMA** and **PUBLIC**. Choose **PUBLIC**

Source Schemata Selection
Choose the schemata you want to migrate.

Schemata

Select all schemata that have to be migrated.

🔍 _____

Show as List ☐

INFORMATIO...   PUBLIC

1 schema selected.

Object Type Selection

Click Next.

Object Type Mapping

Click **Show Details** on both sections. For **Migration Method for Type Schema**, choose **Multilanguage**. For **Migration Method for Type Table**, choose **Data Consistancy/Multilanguage**

Click **Advanced**. Check **Enabled Detailed Mappings in Next Step**

Detailed Object Mapping

Click to rename the **destination database** to be the one set aside to migrate to.

From this point on, you should be able to click next all the way through to finish the migration.

Viewing and Editing License Details

The '**License Details**' page tells you:

- How many users your Confluence instance is licensed to support, and how many are currently registered.
  Note: The number of registered users only includes users who have '**can use Confluence**' permission. Deactivated users are not included.
  Click the '**Refresh**' button to make sure you see the latest count.
- What type of license you have (e.g. Commercial, Academic, Community).
- How much time remains in your one-year support and upgrades period (for full licenses) or 30-day trial (for trial licenses).
- Your server ID, which:
  - is generated when you install Confluence for the first time
  - exists for the life of the Confluence instance
  - survives an upgrade
  - is held in the database
  - is not bound to a specific license
  - is the same for all servers in a cluster.

To view the details of your Confluence license,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.

2. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Click '**License Details**' under the heading '**Administration** in the left-hand panel.

To look up your license details on the Atlassian website, please log in to my.atlassian.com

To update your Confluence license,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.

2. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Click '**License Details**' under the heading '**Administration** in the left-hand panel.
4. Enter your new license details into the '**License**' field and click the '**Save**' button.

   If you are running a Confluence Cluster, you will need to:
   - Update each server's Confluence license separately.
   - Ensure that the new license has enough nodes to cover all servers that are currently running in your cluster. (To check the number of active servers in your cluster, see the Cluster Administration page.)

Screenshot : License Details

This page shows your current licensing information.

You can use the form below to update the license Confluence is running with.

Organisation	**Atlassian**
Date Purchased	**Feb 11, 2007**
License Type	**Confluence: Commercial Server**
Licensed Users	**500** (0 signed up currently) [Refresh]
Support Period	Your commercial Confluence support and updates are available until **Feb 12, 2008**.
Server ID	**AACK-COI5-AACK-COI5** (Atlassian sales or support may ask you to provide this ID)
License	[ ] [Save]

Downgrading your Confluence license (reducing the number of allowed users)
If you need to downgrade your Confluence license to one which allows fewer users, please make sure first that your new license covers your current user base.

- View your license details as described above.
- Verify that the number of users '**signed up currently**' is lower than the number allowed by the new license.
- If you currently have more users signed up than the new license allows, please follow these instructions on removing users from your Confluence site.

RELATED TOPICS

Viewing and Editing License Details

Cache Statistics

Administrators Guide Home   Confluence Documentation Home

Viewing System Information

The System Information screen provides information about Confluence's configuration, and the environment in which Confluence has been deployed. Your system configuration information is helpful to us when diagnosing errors you may face using Confluence. If you file a support request or bug report, the more detail you can provide about your installation and environment the faster we will be able to help.

To view your system information,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Configuration**' in the 'Administration' section.

> ✅   The handy **Memory Graph** helps you keep track of Confluence's memory usage.

RELATED TOPICS

- Cache Statistics
- Site Statistics
- Viewing and Editing License Details
- Viewing and Managing Installed Plugins
- Live Monitoring Using the JMX Interface

Administrators Guide Home   Confluence Documentation Home

Live Monitoring Using the JMX Interface

> ⚠   Like the Heisenberg uncertainty principle, adding live monitoring to a production instance may have an impact itself on performance!

With the JMX interface (introduced in Confluence 2.8), you can monitor the status of your Confluence instance in real time. This will provide you with useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues. To read the JMX data, you will need to use a JMX client.

Disable JMX

> ⚠   If you experience any problems during Confluence startup that are related to JMX, it is possible to disable the JMX registration process. Please place jmxContext.xml in your `<confluence-install>/confluence/WEB-INF/classes` folder to do so.

What is JMX?

JMX (Java Management eXtensions) is a technology for monitoring and managing Java applications. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application.

1. Enabling JMX Remote with Tomcat

By default, Confluence uses the Apache Tomcat web server. To use JMX, you must enable it on your Tomcat server, by carrying out the steps under the Apache Tomcat documentation, entitled Enabling JMX Remote. With those steps completed, restart your Tomcat server.

For the stand-alone, add the startup parameter -Dcom.sun.management.jmxremote to setenv.sh or setenv.bat. See instructions for the Windows Service - enter it in the same place as PermGen Memory.

2. Selecting your JMX Client

You need to use a JMX client in order to view the JMX output from Confluence. JConsole is a readily available JMX client that is included with Sun's Java Developer Kit (version 5 onwards). The full name is the '*Java Monitoring and Management Console*', but we will refer to it as JConsole for the purposes of this document.

3. Adding the JMX Client to your Path

You must add the location of the JConsole binary file to your '**path**' environment variable. As JConsole resides in the 'bin' (binaries) folder under your Java directory, the path should resemble something like this:

```
1. JDK_HOME/bin/
```

In this example, replace 'JDK_HOME' with the full system path to your Java directory.

4. Configuring JConsole

To configure JConsole,

1. Run the JConsole application.

2. You will be prompted to create a new connection. Choose '**remote process**', enter the hostname of your Confluence instance and a port of your choosing.

   > ✔ To connect easily, add the startup parameters to setenv.bat or setenv.sh:
   > -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8086
   > -Dcom.sun.management.jmxremote.authenticate=false
   > Port 8086 is unlikely to be used. Then, connect remotely using port 8086.

3. Click '**Connect**'.

Note: Other JMX clients besides JConsole can read JMX information from Confluence.

What can I monitor with JMX?

The JMX interface allows you to see live internal information from your Confluence instance, via the following MBeans:

IndexingStatistics

This MBean shows information related to search indexing.

Property name	Function	Values
Flushing	Shows state of cache (i.e. flushing, or not).	True/False
LastElapsedMilliseconds	Time taken during last indexing.	Milliseconds
LastElapsedReindexing	Time taken during last re-indexing.	Milliseconds
TaskQueueLength	Shows number of tasks in the queue.	Integer

SystemInformation

This MBean shows information related to database latency. It also contains most of the information presented on the System Information page.

Property name	Function	Values
DatabaseExampleLatency	Shows the latency of an example query performed against the database.	Milliseconds

RequestMetrics

This MBean shows information related to system load and error pages served.

Property name	Function	Values

AverageExecutionTimeForLastTenRequests	Average execution time for the last ten requests.	Milliseconds
CurrentNumberOfRequestsBeingServed	Number of requests being served at this instant.	Integer
ErrorCount	Number of times the Confluence error page was served.	Integer

MailServer-SMTPServer

This MBean shows information related to email dispatch attempts and failures. There will be an MBean for every SMTP Mailserver that has been configured in the Confluence instance.

Property name	Function	Values
EmailsAttempted	The number of email messages Confluence has tried to send.	Integer
EmailsSent	The number of email messages sent successfully.	Integer

MailTaskQueue

This MBean shows information related to the email workload.

Property name	Function	Values
ErrorQueueSize	Number of errors in the queue.	Integer
Flushing	Shows state (i.e. flushing, or not)	True/False
FlushStarted	Time that operation began.	Time
RetryCount	The number of retries that were performed.	Integer
TaskSize	Number of email messages queued for dispatch.	Integer

SchedulingStatistics

This MBean shows information related to current jobs, scheduled tasks and the time that they were last run.

High CPU consuming threads

For Java 1.6, add the Top Threads Plugin to monitor whether CPU is spiking. Download it to a directory and run JConsole like this: JConsole -pluginpath /pathto/topthreads.jar

This works only with jdk 1.6, but that can be on the remote machine if the server is running a lower version.

RELATED TOPICS

- Viewing System Information
- Cache Statistics
- Viewing and Editing License Details
- Viewing and Managing Installed Plugins

Site Statistics

Understanding Site Statistics

Site Statistics allows you to view a breakdown of pages and editing activity on your Confluence instance.

You can also choose to email your site statistics information to Atlassian, with the '**Mail Statistics**' button. This may be useful for troubleshooting purposes, or when requesting Atlassian support. You can also review the data before it is sent (if you are concerned that it may contain sensitive information).

Screenshot: Confluence Site Statistics

Viewing Confluence Site Statistics

To view Confluence Site Statistics,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.

2. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**Site Statistics**' under the heading 'Administration' in the left-hand panel.
4. Click '**Load Statistics**'. For a large Confluence installation, this may affect the server's performance. For this reason, we suggest you do this at a 'quiet time' for the system.
5. Site Statistics will be displayed.

🚫 The Global Statistics Plugin is unsupported.

⚠️ If the '**Site Statistics**' option is not shown, you may have to install the **Global Statistics Plugin**. In this case, follow the instructions for installing plugins and look for the '**Global Statistics Plugin**'.

✅ If you want to view page-view statistics, this is possible with the SQL Plugin. Some examples are available on the Confluence Community Pages. Please note that some plugins are not officially supported by Atlassian.

Installing Patched Class Files

Atlassian support or the Atlassian bug-fixing team may occasionally provide patches for critical issues that have been resolved but have not yet made it into a release. Those patches will be class files which are attached to the relevant issue in our JIRA bug-tracking system.

Installation Instructions for Confluence Standalone

Follow these steps to install a patched class file:

1. Shut down your confluence instance.
2. Copy the supplied class files to `<installation-directory>/confluence/WEB-INF/classes/<subdirectories>`, where:
   - `<installation-directory>` must be replaced with your Confluence Installation directory. (If you need more information, read about the Confluence Installation Directory.)
   - `<subdirectories>` must be replaced by the value specified in the relevant JIRA issue. This value will be different for different issues. In some cases, the subdirectories will not exist and you will need to create them before copying the class files. Some issues will contain the patch in the form of a ZIP file which will contain the desired directory structure.
3. Restart your Confluence instance for the changes to become effective.

Class files in the `/WEB-INF/classes` directory of a web application will be loaded before classes located in JAR files in the `/WEB-INF/lib` directory. Therefore, classes in the first directory will effectively replace classes of the same name and package which would otherwise be loaded from the JAR files.

RELATED TOPICS

Editing Files within JAR Archives
Where are the files that used to be in my Confluence installation directory?
Administrators Guide Home Confluence Documentation Home

Configuring Confluence

- Site Configuration
  - Configuring the Server Base URL
  - Configuring the Site Homepage
  - Configuring the Site Support Address
  - Customising Default Space Content
  - Editing the Global Logo
  - Editing the Site Title
  - Editing the Site Welcome Message
  - Showing Link Icons
  - View Space Goes to Browse Space
- Configuring Encoding
  - Character encodings in Confluence
  - Troubleshooting Character Encodings
    - "€" Euro character
    - MySQL 3.x Character Encoding Problems
- Configuring Mail
  - Configuring a Server for Outgoing Mail
  - Enabling the 'Mail Page' plugin
  - The Mail Queue
- Optional Settings
  - Attachment Storage Configuration
    - Hierarchical File System Attachment Storage
  - Configuring Quick Navigation
  - Enabling CamelCase Linking
  - Enabling OpenSearch
  - Enabling Remote APIs
  - Enabling Rich Text Editing Option
  - Enabling the Did You Mean Feature
  - Enabling Threaded Comments
  - Enabling Trackback
  - Making Rich Text Editing default
  - WebDAV Configuration
- Other Settings
  - Configuring Attachment Size
  - Configuring Character Encoding
  - Configuring HTTP Timeout Settings
  - Configuring Indexing Language
  - Configuring Jira Issues Icon mappings
  - Configuring Number Formats
  - Configuring Shortcut Links
  - Configuring Time and Date Formats
  - Number of Ancestors to Show in Breadcrumbs
  - Thumbnail Settings
- Configuring System Properties
  - Recognised System Properties
- Configuring Logging
- External Gadgets

Administrators Guide Home Confluence Documentation Home

Site Configuration

- Configuring the Server Base URL
- Configuring the Site Homepage
- Configuring the Site Support Address
- Customising Default Space Content
- Editing the Global Logo
- Editing the Site Title
- Editing the Site Welcome Message
- Showing Link Icons
- View Space Goes to Browse Space

Configuring the Server Base URL

The **Server Base URL** is the URL via which users access Confluence. The base URL **must** be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

ℹ️ You need to have System Administrator permissions in order to perform this function.

To configure the Server Base URL,

1. In Confluence, open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' will open.
2. Click '**General Configuration**' in the left-hand panel.
3. Click the '**Edit**' button next to '**Site Configuration**'.
4. Enter the new URL in the '**Server Base URL**' text box.
5. '**Save**' your changes.

⚠️ If you configure a different base URL or if visitors use some other URL to access Confluence, it is possible that you may encounter errors while viewing some pages.

ℹ️ **Example**

If Confluence is installed to run in a non-root context path (that is, it has a context path), then the server base URL should include this context path. For example, if Confluence is running at `http://www.foobar.com/confluence`, the server base URL should be `http://www.foobar.com/confluence`.

RELATED TOPICS

📄 Customising Default Space Content

📄 Editing the Site Welcome Message

📄 View Space Goes to Browse Space

📄 Configuring the Site Support Address

📄 Configuring the Site Homepage

📄 Editing the Site Title

📄 Configuring the Server Base URL

📄 Editing the Global Logo

📄 Showing Link Icons

Configuring the Site Homepage

You can configure Confluence to direct users to any of the space home pages on the site when they log in, rather than to the Dashboard.

To configure the site-wide home page,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.

2. Click '**Edit**' next to the '**Site Configuration**' panel.

3. Select a space from the '**Site Homepage**' dropdown menu. When users log in, Confluence will open the home page of the space you choose here.

4. Ensure that the '**View Space Goes to Browse Space**' option is set to '**Off**' if you want users to be sent to the space **home page** and not the **space summary page**.

5. Click the '**Save**' button at the bottom of the screen.

The spaces available to be set as your home page will depend on the access permissions of the space and the site.

- If your site allows anonymous access, the site home page must also be anonymously accessible.
- The site home page must be accessible to the '**confluence-users**' group.

Screenshot : Configuring the site homepage

RELATED TOPICS

📄 Customising Default Space Content

📄 Editing the Site Welcome Message

📄 View Space Goes to Browse Space

📄 Configuring the Site Support Address

📄 Configuring the Site Homepage

📄 Editing the Site Title

📄 Configuring the Server Base URL

📄 Editing the Global Logo

📄 Showing Link Icons

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Configuring the Site Support Address

The **Site Support Address** is an email address which points to a JIRA instance configured to receive and handle support requests by email.

By default, the site support address is set to the Atlassian Support System. In most cases, there is no need to change it.

In order to use the site support address, ensure that SMTP email is set up on your Confluence instance.

To configure the site support address,

1. Go to the **Administration Console** and click 'General Configuration' in the left-hand panel.
2. Click the '**Edit**' button next to 'Site Configuration'.
3. Enter the new '**Site Support Address**'.
4. Click the '**Save**' button at the bottom of the screen.

RELATED TOPICS

Troubleshooting Problems & Requesting Technical Support
Site Configuration
Administrators Guide Home   Confluence Documentation Home

Customising Default Space Content

Confluence Administrators can define **default content for a space home page**. This content will appear on the home page whenever someone adds a new space. You can define different content for global spaces and for personal spaces.

The default content will appear only for new spaces created after you have defined the content. Content in existing home pages will not be changed.

To define default content for home pages in global spaces,

1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left-hand panel.
2. The '**Space Home Pages**' tab will open on the 'Default Space Content' page. Enter the content which you want to appear on the home page for new global spaces. You can use special characters within the content as variables (place holders). Confluence will replace the curly brackets and digits with the corresponding information as shown below:
    - **{0}** — The space name.
3. Click the '**Save**' button.

To define default content for home pages in personal spaces,

1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left panel.
2. The 'Space Home Pages' tab will open on the 'Default Space Content' page. Click the '**Personal Space Home Pages**' tab.
3. Enter the content which you want to appear on the home page for new personal spaces. You can use special characters within the content as variables (place holders). Confluence will replace the curly brackets and digits with the corresponding information as shown below:
    - **{0}** — The space owner's full name.
    - **{1}** — The space owner's e-mail address.
    - **{2}** — Any personal information the space owner has entered on their user profile in the 'Information about me' section.
4. Click the '**Save**' button.

You can also undo all customisations of the default home page content, and go back to the default content as originally supplied with Confluence.

To restore the original default content,

1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left panel.
2. Select either the 'Space Home Pages' tab or the '**Personal Space Home Pages**' tab, as required.
3. Click the '**Revert**' button.

Screenshot : Defining default space content

Default Space Content

Configuration

- General Configuration
- Daily Backup Admin
- Manage Referrers
- Plugins
- Plugin Repository
- Languages
- Shortcut Links
- Global Templates
- Mail Servers
- User Macros
- JIRA Issues Icon Mappings
- Attachment Storage
- Spam Prevention
- PDF Language Support
- Default Space Content
- Configure Whitelist
- Office Connector Configuration

Look and Feel

- Themes
- Colour Scheme
- Layouts
- Stylesheet
- Global Logo

Default Space Content

This page allows you to customise what is shown on the *home page* of a space when it is first created. Customisations will only apply to spaces created after you have made your changes. Existing spaces will remain unchanged.

Space Home Pages | **Personal Space** Home Pages

```
This is the home of the {0} space.

To help you on your way, we've inserted some of our favourite macros on this home page. As
you start creating pages, adding news items and commenting you'll see the macros below fill
up with all the activity in your space.
{section}
{column:width=60%}
{recently-updated}
{column}
{column:width=5%}
{column}
{column:width=35%}
h6. Navigate space
{pagetreesearch}
{pagetree}
{column}
{section}
```

[Revert] [Save]

The symbol **{0}** will be replaced with the name of the Space.

RELATED TOPICS

- Customising Default Space Content
- Editing the Site Welcome Message
- View Space Goes to Browse Space
- Configuring the Site Support Address
- Configuring the Site Homepage
- Editing the Site Title
- Configuring the Server Base URL
- Editing the Global Logo
- Showing Link Icons

Administrators Guide Home   Confluence Documentation Home

Editing the Global Logo

By default, the global logo appears beside the page title on all pages in the site. You can disable the logo or replace it with one of your own.

To edit the global logo,

1. From the '**Administration Console**' click on '**Global logo**' under the heading '**Look and Feel**' in the left panel.

2. In the screen displayed, select '**Off**' to disable logo.

3. To upload a new logo, click '**Browse**' to select a new image and click '**Upload Logo**'.

RELATED TOPICS

Customising Default Space Content

Editing the Site Welcome Message

View Space Goes to Browse Space

Configuring the Site Support Address

Configuring the Site Homepage

Editing the Site Title

Configuring the Server Base URL

Editing the Global Logo

Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Editing the Site Title

The site title appears in your browser's title bar. By default, it is set to 'Confluence'.

To change the title of your Confluence instance,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.

3. Enter a new title for your site in the input field beside '**Site Title**' and '**Save**'.

RELATED TOPICS

Customising Default Space Content

Editing the Site Welcome Message

View Space Goes to Browse Space

Configuring the Site Support Address

Configuring the Site Homepage

Editing the Site Title

Configuring the Server Base URL

Editing the Global Logo

Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Editing the Site Welcome Message

The site welcome message appears on the Dashboard. It can be used to provide users with an introduction to the site, or as a "message of the day".

To edit the site welcome message,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the top of the '**Site Configuration**' section.

3. In the text-entry box beside '**Site Welcome Message**' enter your text using regular Confluence markup.

4. '**Save**' your changes.

How we use it at Atlassian

Atlassian makes great use of the welcome message feature on our internal Confluence wiki. You can see it here:

(click to expand)

The welcome message itself is just an include macro:

```
1.{include:STAFF:Extranet Homepage}
```

The include macro allows an entire page to be shown in line. This particular page lives in the STAFF space, where anyone can edit it. We usually have some amusing picture, or company-wide notice showing on it. The featured photo generally changes each week – you can see the current (29 June 2009) homepage to the right. The page itself has over 600 edits by many different people.

The page also includes an edit link, for quick access to change the welcome message.

Our homepage wiki markup looks something like this:

```
01.!Clover Dukey.jpg|width=200!
02.
03.{nodisplay}
04.This is the content that goes on the Extranet homepage, above the spaces list.
05.
06.NOTE: KEEP YOUR PICTURES SMALL (<80KB) -- USE JPG FOR PICTURES, WIDTH 400
07.{nodisplay}
08.h4. Experimental blogroll: All posts labelled "extranet-dashboard"
09.
10.{blog-posts:content=titles|labels=extranet-dashboard|spaces=@all|max=10}
11.If you want to promote a good post to stand out from the eac white noise,
12.just add the label *extranet-dashboard*. To avoid inflation please use the
13.label carefully.
14.
15.{float-right}
16.([edit me|http://extranet.atlassian.com/pages/editpage.action?pageId=603422736])
17.{float-right}
```

RELATED TOPICS

📄 Customising Default Space Content

📄 Editing the Site Welcome Message

📄 View Space Goes to Browse Space

📄 Configuring the Site Support Address

📄 Configuring the Site Homepage

📄 Editing the Site Title

📄 Configuring the Server Base URL

📄 Editing the Global Logo

📄 Showing Link Icons

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Showing Link Icons

In order to distinguish external links, user links and email links in wiki content, the Confluence Administrator can configure Confluence to show a small icon in the top right-hand corner of each link.

http://www.google.com⬈

Company Documentation✚

sales@atlassian.com✉

To show link icons,

1. From the '**Administration Console**', click '**General Configuration**' under the heading '**Configuration**' in the left-hand panel.

2. Click the '**Edit**' button next to '**Formatting and International Settings**'.

3. Beside '**Show Link Icons**', select '**On**' to enable the feature. Select '**Off**' to disable it.

4. Click '**Save**'.

RELATED TOPICS

📄 Customising Default Space Content

📄 Editing the Site Welcome Message

📄 View Space Goes to Browse Space

📄 Configuring the Site Support Address

📄 Configuring the Site Homepage

📄 Editing the Site Title

📄 Configuring the Server Base URL

📄 Editing the Global Logo

📄 Showing Link Icons

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

View Space Goes to Browse Space

By default, when you click on a space link, you are taken to the space's homepage. If you wish, you can configure Confluence to redirect all space links on the site to the '**Browse Space**' view of the space instead.

To direct the space link to the 'browse space' view,

1. Go to the '**Administration Console**' click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen. .

3. Select '**ON**' beside '**View Space goes to Browse Space**' and click '**Save**'.

RELATED TOPICS

Customising Default Space Content

Editing the Site Welcome Message

View Space Goes to Browse Space

Configuring the Site Support Address

Configuring the Site Homepage

Editing the Site Title

Configuring the Server Base URL

Editing the Global Logo

Showing Link Icons

Administrators Guide Home   Confluence Documentation Home

Configuring Encoding

Confluence allows the configuration of which character encoding is used to deliver pages.

> ⚠ While different character encodings are supported, we strongly recommend that **UTF-8** is used. Confluence is heavily tested on UTF-8, and users are likely to have less problems with this encoding than others.

> ⚠ **Mac Users**
> Mac Users please note that **MacRoman** encoding is compatible with UTF-8. You do not need to change your encoding settings if you are already using MacRoman.

To avoid problems with character encoding, make sure the encoding used across the different components of your system are the same:

- Configuring Database Character Encoding
- Application Server URL encoding
- Confluence Character Encoding

If you are having problems with the character encoding in Confluence, please see the Troubleshooting Character Encodings page.

Character encodings in Confluence

Where character encoding is used

There are three places that character encoding matters to Confluence:

1. **Database encoding** - usually the most important; it is where almost all user data is stored.
2. **Filesystem encoding** - important for attachment storage (pre-2.2), reading Velocity templates and writing exported files.
3. **HTTP request and response encoding** - important for form parsing, correct rendering by the browser and browser interpretation of encoded URLs.

Problems generally arise when Confluence thinks one of the above encoding is different to what it actually is. For example, Confluence might believe the database is using ISO-8859-1 encoding, when in fact it is UTF-8 encoded.

Java character encoding

Java *always* uses the double-byte UCS-2 character encoding for all `char` and `String` data. This means that each of the encodings above defines how, at that particular point, characters are converted to and from Java's native UCS-2 format into some other format that the browser, filesystem or database might understand.

So when a request comes in to Confluence, we convert it from the request encoding to UCS-2. Then we store that data into the database, converting from UCS-2 to the database's encoding. Retrieving information from the database and sending it back to the browser is the same process in the opposite direction.

Problems with character encodings

If Confluence has the wrong idea about encoding for one of the above, it manifests itself in different ways:

1. Incorrect database encoding - user data is corrupted between saving and restoring from the database. This often happens after a delay, as we cache data as it is written to the database and only later retrieve the corrupted copy from the database.
2. Incorrect/non-Unicode filesystem encoding - international filenames break attachment download/upload/removal (pre-2.2); exports break with international content or attachments.
3. Incorrect HTTP encoding - incorrect encoding selected by browser, resulting in incorrect rendering of characters. Changing browser encoding causes page to render properly. Broken URLs when linking to pages or attachments with non-ASCII characters.

Configuration of character encodings

The **Confluence character encoding** is a configuration setting found in `Administration > General Configuration`, and at runtime available in Settings.defaultEncoding. It is subsequently used in the following parts of the system:

- ConfluenceWebWorkConfiguration sets `webwork.i18n.encoding` to the this encoding, which WebWork uses in the response Content-Type header.
- AbstractEncodingFilter sets the HTTP request encoding to this encoding. This seems unnecessary, since the Content-Type header from the client should include the encoding used. This affects form submissions and file uploads.
- VelocityUtils reads in Velocity templates using this encoding when reading templates from disk.
- AbstractXmlExporter creates its output using this encoding.
- GeneralUtil uses this encoding when doing URLEncode and URLDecode. Different browsers have different support for character sets in URLs, so it's uncertain how much benefit this provides.

In summary, changing the Confluence character encoding will change your **HTTP request and response encoding** and your **Filesystem encoding** as used by exports and velocity templates.

The **database encoding** is the responsibility of your JDBC drivers. The drivers are responsible for reading and writing from the database in its native encoding and translating this data to and from Java Strings (which are UCS-2). For some drivers, such as MySQL, you must set Unicode encoding explicitly in the JDBC URL. For others, the driver is smart enough to determine the database encoding automatically.

Ideally, your database itself should be in a Unicode encoding (and we recommend doing this for the simplest configuration), but that is not necessary as long as:

- the database encoding supports all the characters you want to store in Confluence
- your JDBC drivers can properly convert from the database encoding to UCS-2 and vice-versa.

The **filesystem encoding** is mostly ignored by Confluence, except for the cases where the above configuration setting above plays a part (exports, velocity). When attachments are uploaded, they are written as a stream of bytes directly to the filesystem. It is the same when they are downloaded: the bytes from the file InputStream are written directly to the HTTP response.

In some places in Confluence, we use the *default filesystem encoding* as determined by the JVM and stored in the `file.encoding` system property (it can be overridden by setting this property at startup). This encoding is used by the Java InputStreamReader and InputStreamWriter classes by default. This encoding should probably never be used; for consistent results across all filesystem access we should be using the encoding set in the General Configuration.

In certain cases we explicitly hard-code the encoding used to read or write data to the filesystem. Two important examples are:

- importing Mbox mailboxes which are known to be ISO-8859-1
- Confluence Bandana config files are always stored as UTF-8.

Some application servers, Tomcat for example, have an encoding setting that modifies Confluence URLs before they reach the application. This can prevent access to international pages and attachments (really anything with international characters in the URL). See configuring your Application Server URL encoding.

Advice

In general, **always set all character encodings to UTF-8**. That includes database, JDBC drivers, application server, filesystem and Confluence.

In certain isolated cases (e.g. Microsoft Windows), it might not be possible to use a fully Unicode filesystem (that is, a default Windows install doesn't support Unicode filenames properly). If so, stick with UTF-8 for the other two and be aware that your operating system might have limitations around international attachments (pre-2.2), backup and restore of international data, etc.

RELATED TOPICS:

- Configuring Database Character Encoding
- Troubleshooting Character Encodings

Troubleshooting Character Encodings

Often users may have problems with certain characters in a Confluence instance. Symptoms may include:

- Non-ASCII characters appearing as question marks (?)
- Page links with non-ASCII characters not working
- Single characters being displayed as two characters
- Garbled text appearing

In most cases, it is due to a mis-configuration in one of the components that Confluence uses.

Follow these steps to diagnose the problem:

1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration.

To perform the test, access the Encoding Test page via the `<confluence base-url>/admin/encodingtest.action` page on your Confluence instance. You will be required to copy and paste a line of text and submit a form. The test will take the text and pass it through Confluence, the application server and the database, and return the results.

You should also test pasting some sample text (Japanese for example) if you are experiencing problems with a specific language.

Example:

```
http://confluence.atlassian.com/admin/encodingtest.action
```

or

```
http://<host address>:<port>/admin/encodingtest.action
```

> 🚫 If the text displayed in the encoding test is different to what was entered, then there are problems with your character encoding settings.

A successful test looks like the following:

The encoding test has now been run. Below, you can compare the raw text delivered from Confl a round-trip through the database. All the test results should appear identical.

Iñtërnâtiônàlizætiøn	This image is how all of the test results below *should* appear. this page, and all of your <u>System Information</u>.

Test 1: Raw text

This is the test string generated in Confluence

Iñtërnâtiônàlizætiøn

Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

Iñtërnâtiônàlizætiøn

Test 3: Database round-trip (select as LOWER)

This is the string from Test 2 after being stored in the database and then retrieved

iñtërnâtiônàlizætiøn

Expected result (converting Java string to lowercase)

iñtërnâtiônàlizætiøn

Test 4: Database round-trip (select as UPPER)

This is the string from Test 2 after being stored in the database and then retrieved

IÑTËRNÂTIÔNÀLIZÆTIØN

Expected result (converting Java string to uppercase)

IÑTËRNÂTIÔNÀLIZÆTIØN

> ⚠️ **MySQL 3.x**
> MySQL 3.x is known to have some problems with the upper- and lower-casing of some characters, and may fail the last two tests. For more information, see MySQL 3.x Character Encoding Problems.

2. Ensure the same encoding is used across all components

As mentioned in the Configuring Encoding document, the same character encoding should be used across the database, application server and web application (Confluence).

- To change the character encoding used in **Confluence**, see Configuring Character Encoding.
- To change the character encoding used in the **application server**, please ensure you set the Application Server URL encoding and view your application server's documentation on any other settings required to enable your encoding.
- To change the character encoding used in the **database**, see Configuring Database Character Encoding.

3. Requesting support

If there are still problems with character encoding after following the above steps, create a support request, and our support staff will aid in solving your problem.

Entering in the following details will help us to identify your problem:

- Attach screenshots of the problem
- Attach the results of the encoding test (above)
- Select which application server (and version) you are using
- Select which database (and version) you are using
- Copy the contents of the System Information page into the 'Description' field

"€" Euro character

The € (euro) symbol is a three byte character, with byte values in file (UTF-8) of 0xE2, 0x82, 0xAC.

Sometimes, if the character encoding is not set consistently among all participating entities of the system, Confluence, server and the database, one may experience strange behaviour.

> ...
> I write a page with a Euro sign in it (€). All is well, the Euro sign shows up in the wiki markup text-box, and the preview, and the display of the saved page.
> One day later, the Euro sign has changed into a question mark upside down!
> ...
> What is going on? Why does the Euro sign mysteriously change? How do I prevent it?

Interestingly enough the character encoding test passes with no problems, demonstrating that Confluence and the connected Database both recognise the € symbol.

There are two potential reasons for this behaviour:

Database and Confluence is using utf-8 encoding. The connection is not.

When data transferred to it via the connection which does not use utf-8 encoding gets encoded incorrectly. Hence, updating the connection encoding may resolve this problem from now on, yet it probably would not affect already existing data.

Database is not using utf-8. Confluence and your connection are.

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "€" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

> For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8, the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *latin1* different characters, instead of one utf8 character. Nevertheless, this is not the case for 3-byte utf8 characters, such as the Euro symbol.

Please ensure that you set the character encoding to UTF-8 for all the entities of your system as advised in this guide.

MySQL 3.x Character Encoding Problems

MySQL 3.x is known to have some problems upper- and lower-casing certain (non-ASCII) characters.

Diagnosing the problem

1. Follow the instructions for Troubleshooting Character Encodings.
2. If the upper- and lower-cased strings displayed on the Encoding Test are different, then your database is probably affected.

An example (faulty) output of the Encoding Test is shown below:

Screenshot: Encoding Test Output

The encoding test has now been run. Below, you can compare the raw text delivered from Cc
round-trip through the database. All the test results should appear identical.

Iñtërnâtiônàlizætiøn	This image is how all of the test results below *should* appe this page, and all of your System Information.

Test 1: Raw text

This is the test string generated in Confluence

Iñtërnâtiônàlizætiøn

Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

Iñtërnâtiônàlizætiøn

Test 3: Database round-trip (select as LOWER)

This is the string from Test 2 after being stored in the database and then retrieved

iñtërnâtiônàlizætiøn

Expected result (converting Java string to lowercase)

iñtërnâtiônàlizætiøn

Test 4: Database round-trip (select as UPPER)

This is the string from Test 2 after being stored in the database and then retrieved

IñTËRNÂTIÔNÀLIZæTIØN

Expected result (converting Java string to uppercase)

IÑTËRNÂTIÔNÀLIZÆTIØN

Solution

Upgrade to a newer version of MySQL. (4.1 is confirmed to work.)

Configuring Mail

- Configuring a Server for Outgoing Mail
- Enabling the 'Mail Page' plugin
- The Mail Queue

- Customising the eMail Templates

Administrators Guide Home Confluence Documentation Home

Configuring a Server for Outgoing Mail

Configuring your Confluence server to send outgoing mail allows your Confluence users to:

- Receive Daily Reports.
- Send a page via email.

ⓘ You need to have System Administrator permissions in order to perform this function.

To configure Confluence Standalone to send outgoing mail,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Mail Servers**' under '**Configuration**' in the left panel. This will list all currently configured SMTP servers.
3. Click '**Add New SMTP Server**' (or **edit** an existing server).
   - **Name**: By default, this is set to 'SMTP Server'.
   - **From**: Enter the email address that will be displayed in the 'from' field for email messages originating from this server.
   - **Subject Prefix**: Enter a subject prefix, if required.
4. Configuring the Host Address, Username and Password:
   - Manually enter your '**Host Address**', '**Username**' and '**Password**' details in the form fields displayed (recommended).
     OR
   - Specify the '**JNDI**' location of a mail session configured in your application server in the form field displayed.

Troubleshooting

If you experience problems with these configurations, please check that your `<Confluence-Install>/confluence/WEB-INF/lib` contains only one copy of the following JAR files:

1. activation-x.x.x.jar
2. mail-x.x.x.jar

Ideally, these should be:

- activation-1.0.2.jar
- mail-1.3.2.jar (or later)

You will then need to move these into the proper directory:

Standalone distribution: Please move the above three jar files from the `<Confluence-Install>/confluence/WEB-INF/lib` directory to `<confluence-install>/lib` (for Confluence version 2.10 onwards) or `<Confluence-Install>/common/lib` (for earlier product versions) and restart Confluence.

RELATED TOPICS

📄 Enabling the 'Mail Page' plugin

📄 Configuring a Server for Outgoing Mail

📄 The Mail Queue

Enabling the 'Mail Page' plugin

The 'Mail Page' plugin allows anyone with the 'View' space permission to email a Confluence page.

The 'Mail Page' plugin is disabled by default. This is because, when someone emails a Confluence page, they can select from a list of all Confluence users and groups (note, however, that email addresses are not visible), or even mail the page to arbitrary addresses. If you have enabled anonymous access or self-signup, visitors could potentially use this feature to send spam or nuisance email through your Confluence server.

⚠ This plugin only works when the **mail server** is configured.

ⓘ You need to have System Administrator permissions in order to perform this function.

ⓘ Confluence versions 2.4 and later come with the 'Mail Page' plugin preinstalled.

To enable the 'Mail Page' plugin,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' under '**Configuration**' in the left-hand panel.
3. This will list all plugins that are currently installed in your Confluence system. Click '**Mail Page Plugin**'.
4. This will display the '**Mail Page Plugin**' details. To enable the 'Mail Page' plugin, click '**Enable plugin**'.
5. Ensure that both of the following are enabled:
   - '**Mail Page Link**' — displays the 'E-mail' link next to the 'Copy' link on the 'Page Info' screen (see 'E-mailing a page')
   - '**mailpageactions**' — enables the e-mail operation.

RELATED TOPICS

Enabling the 'Mail Page' plugin

Configuring a Server for Outgoing Mail

The Mail Queue

The Mail Queue

Email messages waiting to be sent out are queued in a mail queue and periodically flushed from Confluence once a minute. A Confluence administrator can also manually flush emails from the mail queue.

If there is an error sending messages, the failed emails are sent to an error queue from which you can either try to resend them or delete them.

To view the mail queue,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Mail Queue**' in the left-hand panel. This will display the emails currently in the queue.
3. Click '**Flush Mail Queue**' to send all emails immediately.
4. Click '**Error Queue**' to view failed email messages. You can try to '**Resend**' the messages, which will flush the mails back to the '**Mail Queue**' or '**Delete**' them from here.

RELATED TOPICS

Enabling the 'Mail Page' plugin

Configuring a Server for Outgoing Mail

The Mail Queue

RELATED TOPICS

Administrators Guide Home Confluence Documentation Home

Optional Settings

- Attachment Storage Configuration
- Configuring Quick Navigation
- Enabling CamelCase Linking
- Enabling OpenSearch
- Enabling Remote APIs
- Enabling Rich Text Editing Option
- Enabling the Did You Mean Feature
- Enabling Threaded Comments
- Enabling Trackback
- Making Rich Text Editing default
- WebDAV Configuration

Attachment Storage Configuration

Confluence allows you to store attachments in one of three places:

- Filesystem - locally in the Confluence home directory
- Database - in Confluence's configured database
- WebDAV - remotely on a WebDAV server (***deprecated***)

A System Administrator can configure Confluence's attachment storage via the '**Attachment Storage**' option on the '**Administration Console**'.

You need to have System Administrator permissions in order to perform this function.

Attachment Storage Options

Local File System

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder. If you are looking to run Confluence Clustered, attachments must be stored in the database.

Database

Confluence gives administrators the option to store attachments in the database that Confluence is configured to use.

Here are some reasons why, as an administrator, you may want to choose this storage system:

- Ease of backup.
- Avoiding issues with certain characters in attachment file names.

> ⚠ While storing attachments in the database can offer some advantages, please be aware that the amount of space used by the database will increase because of the greater storage requirements.

WebDAV

Confluence also allows administrators to set an external WebDAV repository as the location for attachment storage.

> ⚠ **WebDAV attachment manager deprecated**
> The option to store Confluence attachments on a WebDAV server has never worked in a useful fashion, and has not been maintained for many versions.
>
> - The WebDAV attachment manager will be **deprecated** from Confluence 2.7, and will be removed from a later version of Confluence.
> - If you store attachments on external WebDAV servers, we recommend that you migrate to file-system or database-backed attachment storage as soon as possible. Refer to CONF-9313 and CONF-2887.
> - This DOES NOT affect the operation of the WebDAV plugin.

Migration between Attachment Storage Systems

You can 'migrate' your attachments from one storage system to another. All existing attachments will be moved over to the new attachment storage system.

> ⛔ When the migration occurs, all other users will be locked out of the Confluence instance. This is to prevent modification of attachments while the migration occurs. Access will be restored as soon as the migration is complete.

> ⚠ When migrating attachments from your database to a filesystem, the attachments are removed from the database after migration. However, when migrating attachments from a filesystem to your database, the attachments remain on the filesystem after migration. If you wish to change this function's behaviour from 'copy' to 'move', please see CONF-14802 and cast your vote.

To perform a migration, follow the steps below:

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.

2. Click '**Attachment Storage**' in the left-hand panel. The current configuration will be displayed.

Attachment storage configuration

3. Click the '**Edit**' button to modify the configuration.
4. Select the storage system you desire.

Edit attachment storage

5. Click the '**Save**' button to save the changes.
6. A screen will appear, asking you to confirm your changes. Clicking 'Migrate' will take you to a screen that displays the progress of the migration.

Migration warning

Troubleshooting

To enable debug logging for WebDAV attachment storage, add the following to the bottom of `WEB-INF/classes/log4j.properties` and restart Confluence:

```
1. log4j.logger.com.atlassian.confluence.pages.persistence.dao=DEBUG,confluencelog
2. log4j.additivity.com.atlassian.confluence.pages.persistence.dao=false
3.
4. log4j.logger.org.apache.webdav=DEBUG,confluencelog
5. log4j.additivity.org.apache.webdav=false
```

RELATED TOPICS

Important Directories and Files

Administrators Guide Home   Confluence Documentation Home

Hierarchical File System Attachment Storage

Introduction

For Confluence version 3.0, the structure of attachments stored on the filesystem was changed. In versions of Confluence prior to 3.0, attachments were stored in directories corresponding to the id of the content to which they belong. The more content in Confluence with attachments, the more directories you would have immediately beneath your configured attachments directory. This directory structure has been changed in Confluence 3.0 and since the default configuration of Confluence is to store attachments in the filesystem, this change is likely to have relevance to administrators of most existing Confluence installations.

If you are installing Confluence for the first time, there will be no consequences as a result of this change. If you are upgrading from a previous version of Confluence, the migration to this new filesystem structure should happen automatically during the upgrade.

The reason for introducing this change was to address the issue CONF-13004. Certain file systems have a limit on the number of files that can be stored in a directory and large Confluence installations were reaching this limit. In addition, storing too many files at a single directory level can cause performance degradation in some circumstances. This new attachment storage strategy ensures this will no longer be the case.

> ⛔ **Backup Confluence Home**
> Before upgrading to Confluence 3.0, as with any upgrade you must ensure you have a backup of your Confluence home directory before you proceed.

The New Directory Layout

The attachment storage layout was chosen to fulfil the following main requirements:

1. Limit the number of entries at any single level in a directory structure.
2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces (see the JIRA issue for more details).

An attachment in Confluence can be thought of as having a number of identifying attributes: *id*, *space id* and *content id*. That is to say, the attachment logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For attachments within a space in Confluence, the directory structure is typically 8 levels, with the name of each directory level based on the following algorithm:

level	Derived From
1 (top)	Always 'ver003' indicating the Confluence version 3 storage format
2	The least significant 3 digits of the *space id*, modulo 250
3	The next 3 least significant digits of the *space id*, modulo 250
4	The full *space id*
5	The least significant 3 digits of the *content id*, modulo 250
6	The next 3 least significant digits of the *content id*, modulo 250
7	The full *content id*
8	The full *attachment id*

Within the 8th level will be a file for each version of that attachment, named to match the version number e.g. 1

An example:

Attachments:

A	B	C	D
id: 745644 space id: 800432 content id: 632780	id: 782234 space id: 800432 content id: 620002	id: 771250 space id: 810032 content id: 603101	id: 701002 *global logo* content id: 511242
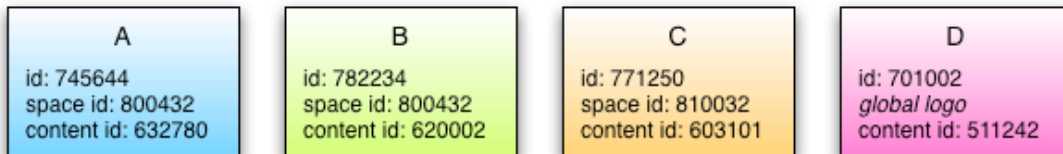
Directory Structure:
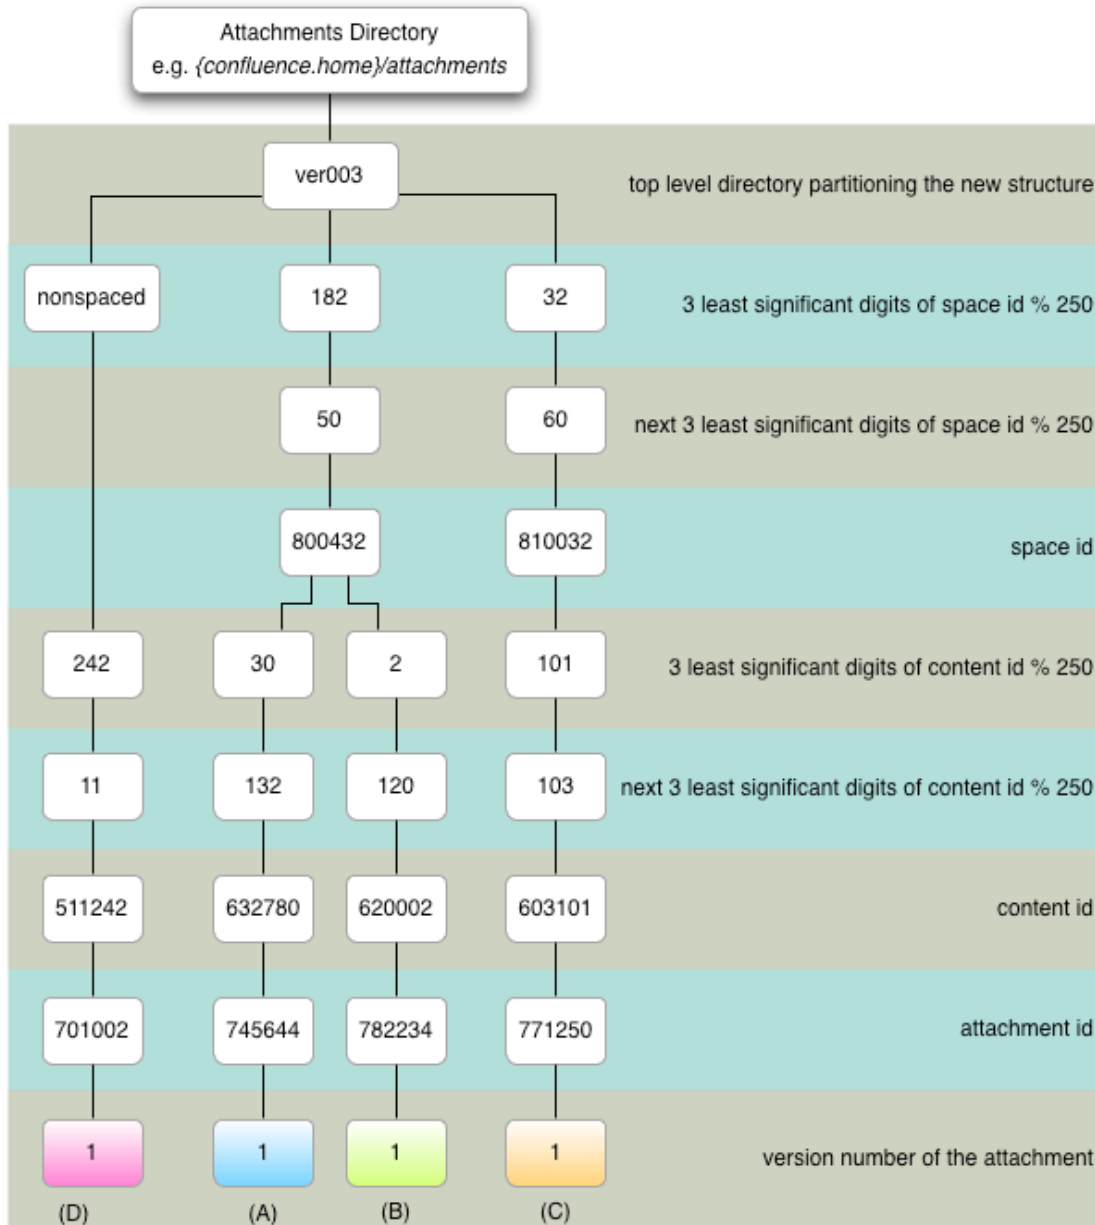
To find the directory where attachments for a particular space are stored, you can use the JSP findspaceattachments.jsp at the location `<confluence url>/admin/findspaceattachments.jsp`. This JSP requires a space key and returns the directory on the file system where attachments for that space are stored.

Attachment D in the above diagram is stored in a slightly different structure. Attachments that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such attachments are the global site logo and also attachments on draft content.

Upgrading to the new attachment storage structure

As mentioned previously, this upgrade is only necessary if you have Confluence configured to store attachments on the file system.

If migration is not necessary due to a different storage configuration (for example, because attachments are stored in the database), then no migration will occur during upgrade and the Confluence log will simply show the following messages -

```
INFO [main] [AbstractUpgradeManager] upgradeStarted Starting automatic upgrade of Confluence
INFO [main] [UpgradeTask] isUpgradeNeeded The configured attachmentDataDao does not store
    attachment data on the file system so the HierarchicalFileSystemAttachmentUpgradeTask is not
necessary.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed successfully
```

Should migration be required, it will occur automatically during upgrade and the log will show output similar to this -

```
INFO [main] [UpgradeTask] doUpgrade Beginning HierarchicalFileSystemAttachmentUpgradeTask.
Depending on the size of the
    attachment data this may take some time.
INFO [main] [UpgradeTask] run 4023 pages may have attachments to be moved to a new hierarchical
structure.
INFO [main] [UpgradeTask] run 0 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 1000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 1500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 2000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 2500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 3000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 3500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 4000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run Successfully moved the attachments for all 4023 pages to the new
hierarchical structure.
INFO [main] [UpgradeTask] doUpgrade Completed HierarchicalFileSystemAttachmentUpgradeTask.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed successfully
```

> ℹ It should be noted that for most implementations of Java, the migration to the new data structure involves moving the files (not copying them). Hence, there should not be a need to have additional disk space available. It also means that the migration should be relatively fast.

Have you previously applied the CONF-8298 patch?

The patch or workaround on the CONF-8298 issue changed the structure of attachment storage but not to the most efficient possible structure. So during the Confluence 3.0 upgrade process this intermediate (CONF-8298) structure will be detected and automatically upgraded.

Troubleshooting the upgrade

⚠ It should be noted that in the event of a failure, your attachment directory may be in an inconsistent state and your first step in troubleshooting should be to restore the backup of your home directory.

There are a number of reasons the migration could fail. This will be shown in the log with a message similar to "`Failed to move the attachments for all pages to the new hierarchical structure.`".

Immediately preceding this message in the log will be entries for each page whose attachments could not be moved. The following table shows examples of these messages and offers some possible explanations.

Example Message	Description
The configured attachment directory `<directory name>` could not be found or was not a Directory.	The configured Confluence attachment directory is not accessible. Check confluence home for the attachment directory and ensure the permissions are correct to allow reading and writing for this directory.
It is not possible to migrate the attachments to the new structure since files already exist which the attachment process may need to create.	Your attachments directory contains files or directories which the upgrade task wants to create. That is, a top level directory called ver003 containing directories or files with names containing up to 3 digits (e.g. 1, 213). This could be due to a previous failed attempt to migrate the attachments. You should restore a previous good copy of your attachments directory and remove any files or directories with this naming pattern before retrying.

Couldn't find current Confluence content for the id `<content Id>`. The attachment is a non-spaced attachment (e.g. global logo, draft attachment, etc) and will be migrated to the nonspaced directory.	This is a normal message indicating that the attachment being migrated does not belong to a space e.g. global logo.
Problem while accessing the database for content id `content Id` so its attachments will not be migrated.	It was not possible to access the database at this point during the migration. You will need restore your Confluence attachment directory from the backup and attempt the upgrade again, once the database is accessible again.
Could not create the new attachment directory `directory`.	The upgrade task could not create the new directory to contain the attachment being moved. Does the server user have sufficient permission to perform this operation in the indicated directory? Is there sufficient disk space?
Failed to move the current attachment directory `<some path>` to the new location of `<some other path>`.	The upgrade task could not move the directory. Does the server user have sufficient permission to perform this operation in the indicated directory?

Configuring Quick Navigation

When a user is searching Confluence (see Using the Quick Navigation Aid) the quick navigation aid automatically offers a dropdown list of pages and other items, matched by title to the search query. By default, this feature is enabled, with the maximum number of simultaneous quick navigation requests set to 40. However, these options can be modified as described below.

> ⓘ The maximum number of simultaneous quick navigation requests defines the maximum number of individuals who can use this feature simultaneously on the same Confluence server. If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To modify the quick navigation feature's options,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. To disable this feature, select '**Off**' beside '**Quick Navigation**'.
5. To modify the maximum number of simultaneous quick navigation requests, enter the appropriate number in the field beside '**Max Simultaneous Requests**'.
6. Click '**Save**'.

The following screenshot demonstrates the user interface of the quick navigation aid.

Screenshot: The quick navigation aid showing titles matching the query 'mark'
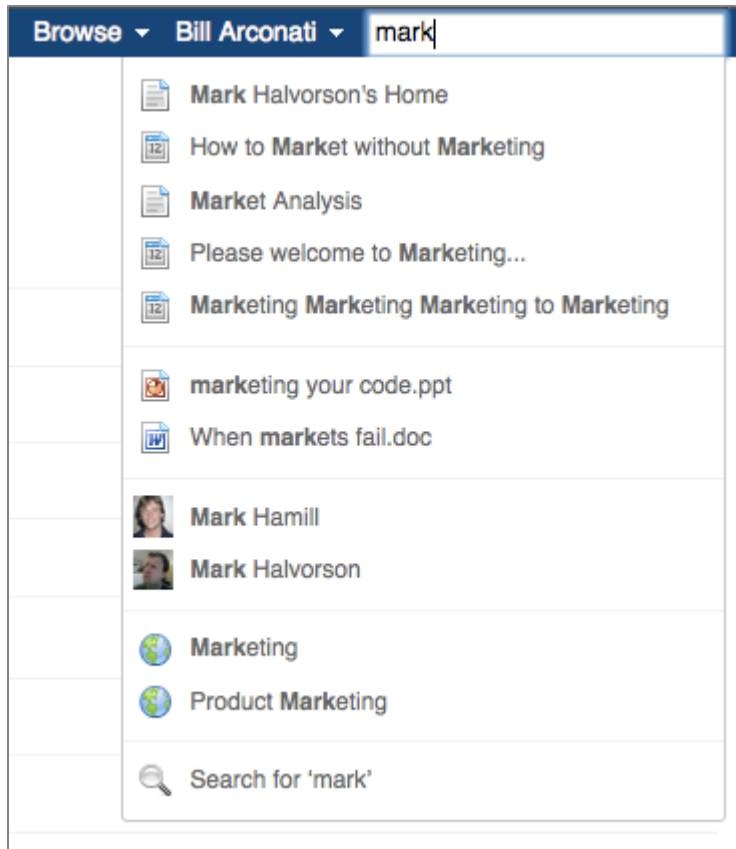
RELATED TOPICS

Searching Confluence

 Administrators Guide Home   Confluence Documentation Home

Enabling CamelCase Linking

CamelCase linking is a form of markup used in many wikis where words are capitalised and compounded together without spaces, 'LikeThis', in order to create links automatically.

By default, CamelCasing is not enabled in Confluence. To use CamelCasing, a Confluence administrator will need to enable this option from the '**Administration Console**'.

To enable CamelCasing,

1. Go to the Confluence '**Administration Console**'. To do this:

    * Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' on the '**General Configuration**' screen.
4. Select '**On**' beside '**CamelCase Links**'.
5. Click '**Save**'.

RELATED TOPICS

 Enabling Remote APIs

 Enabling Trackback

 Enabling Rich Text Editing Option

 Attachment Storage Configuration

 WebDAV Configuration

 Enabling CamelCase Linking

📄 Enabling Threaded Comments

[!Administration Guide Attachments directory^adminhome.gif!] 🏠 Confluence Documentation Home

Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or IE7 search box (see Searching Confluence from your Browser's Search Box). By default, OpenSearch autodiscovery is enabled. This feature can be enabled or disabled as described below.

To enable or disable OpenSearch autodiscovery,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**Open Search**' to enable this feature, or '**Off**' to disable it.
5. Click '**Save**'.

RELATED TOPICS

Searching Confluence

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Enabling Remote APIs

Confluence provides XML-RPC and SOAP remote APIs. You need to enable the APIs from the **Administration Console** before you can access Confluence remotely.

ℹ️ You need to have System Administrator permissions in order to perform this function.

To enable the remote API,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' next to '**Site Configuration**'.
4. Select '**On**' next to '**Remote API (XML-RPC & SOAP)**'.
5. Click '**Save**' to retain your changes.

RELATED TOPICS

No content found for label(s) remoteapi.

Confluencer.NET

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Enabling Rich Text Editing Option

With Confluence 2.0 and later versions, users have the option of using the Rich Text editor to create pages. By default, this is set to 'On'. If desired, a Confluence administrator can disable '**Rich Text Editing**' from the **Administration Console**.

To disable Rich Text editing,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select **General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**Off**' beside '**Rich Text Editing**'.
5. Click '**Save**'.

RELATED TOPICS

📄 Enabling Rich Text Editing Option

📄 Rich Text Editor Overview

📄 Making Rich Text Editing default

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Enabling the Did You Mean Feature

When you perform a full Confluence search, Confluence may offer you an alternative spelling of your search query. The alternative spelling will appear next to the words 'Did you mean'. By default, this feature is disabled and currently only supports the English language. However, it can be enabled as described below.

To enable the 'Did You Mean' feature,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**Did You Mean**'.

    ℹ️ If you have no 'Did you mean' feature index or you have not yet created it, this option will not be available. To create this index, click '**build the did-you-mean index**' and on the subsequent page, click '**Build**' in the '**Did You Mean Index**' section. Then return to the '**General Configuration**' screen in Edit mode.
5. Click '**Save**'.

RELATED TOPICS

Searching Confluence

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Enabling Threaded Comments

Comments on pages or news items are displayed in one of two views:

- **Threaded**: Shows the comments in a hierarchy of responses. Each reply to a comment is indented to indicate the relationships between the comments.
- **Flat**: Displays all the comments in one single list and does not indicate the relationships between comments.

By default, comments are displayed in **threaded** mode. The Confluence administrator can enable or disable the threaded view for the entire Confluence site.

To enable or disable the threaded view,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**Feature Settings**' section, click '**Edit**'.
4. Select '**On**' beside '**Threaded Comments**' to enable threaded mode.
   Select '**Off**' to disable threaded mode and display all comments in flat mode.
5. Click '**Save**'.

RELATED TOPICS

📄 Viewing Comments

📄 Commenting on a Page

Administrators Guide Home   Confluence Documentation Home

Enabling Trackback

When Trackback is enabled, any time you link to an external webpage that supports Trackback Autodiscovery, Confluence will send a trackback ping to that page to inform it that it has been linked to.

Confluence pages also support Trackback Autodiscovery and when Trackback is enabled, can receive trackback pings sent by other sites.

To enable trackback,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. In the '**Feature Settings**' screen, click '**Edit**'.
4. Select "**On**' beside '**Trackback**' and click '**Save**'.

RELATED TOPICS

Enabling or Disabling Public Signup (Confluence Docs 3.1)

Adding SSL for Secure Logins and Page Security (Confluence Docs 3.1)

Hiding external referrers (Confluence Docs 3.1)

Excluding external referrers (Confluence Docs 3.1)

Hiding External Links From Search Engines (Confluence Docs 3.1)

Hiding the People Directory (Confluence Docs 3.1)

User Email Visibility (Confluence Docs 3.1)

Anonymous Access to Remote API (Confluence Docs 3.1)

Ignoring External Referrers (Confluence Docs 3.1)

Spam Prevention via Captcha (Confluence Docs 3.1)

Managing External Referrers (Confluence Docs 3.1)

Administrators Guide Home   Confluence Documentation Home

Making Rich Text Editing default

A Confluence administrator can configure whether the default mode of editing on the site is '**Rich Text**' or '**Wiki Markup**'.

Users will still be able to configure their individual preferences from the '**Edit**' tab of a page.

To make Rich Text Editing the default,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' on the '**General Configuration**' screen.
4. Select '**On**' beside '**Users see Rich Text Editor by default**'.
   (Select '**Off**' to set '**Wiki Markup**' editing as the default.)
5. Click '**Save**'.

RELATED TOPICS

Enabling Rich Text Editing Option

Rich Text Editor Overview

📄 Making Rich Text Editing default

[!Administration Guide Attachments directory^adminhome.gif!] 🏠Confluence Documentation Home

WebDAV Configuration

On this page:

- Introduction to Confluence's WebDAV Client Integration
- Restricting WebDAV Client Write Access to Confluence
- Disabling Strict Path Checking
- Virtual Files and Folders

Introduction to Confluence's WebDAV Client Integration

WebDAV allows users to access Confluence content via a WebDAV client, such as 'My Network Places' in Microsoft Windows. Provided that the user has permission, they will be able to read and write to spaces, pages and attachments in Confluence. Users will be asked to log in and the standard Confluence content access permissions will apply to the equivalent content available through the WebDAV client.

By default, all WebDAV clients have permission to write to Confluence. Write permissions include the ability for a WebDAV client to create, edit, move or delete content associated with spaces, pages and attachments in a Confluence installation.

On the '**WebDAV Configuration**' page, you can:

- Deny a WebDAV client write permissions to a Confluence installation using a regular expression (regex).
- Disable or enable strict path checking.
- Enable or disable access to specific virtual files/folders.

> ℹ️
> - The 'WebDav Configuration' page is only be available if the WebDAV plugin has been enabled. (Refer to Installing and Configuring Plugins using the Plugin Repository Client for more information on enabling Confluence plugins). Note that this plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator.
> - The settings on the 'WebDav Configuration' page do **not** apply to external attachment storage configuration.

Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.

WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:

```
01.PROPFIND /plugins/servlet/confluence/default HTTP/1.1
02.Content-Language: en-us
03.Accept-Language: en-us
04.Content-Type: text/xml
05.Translate: f
06.Depth: 1
07.Content-Length: 489
08.User-Agent: Microsoft Data Access Internet Publishing Provider DAV
09.Host: 127.0.0.1:8082
10.Connection: Keep-Alive
```

> ℹ️ Unlike earlier versions of the WebDAV plugin which could only restrict write permissions for **all** WebDAV clients, the current version of this plugin allows you to restrict write permissions to specific WebDAV clients selectively.
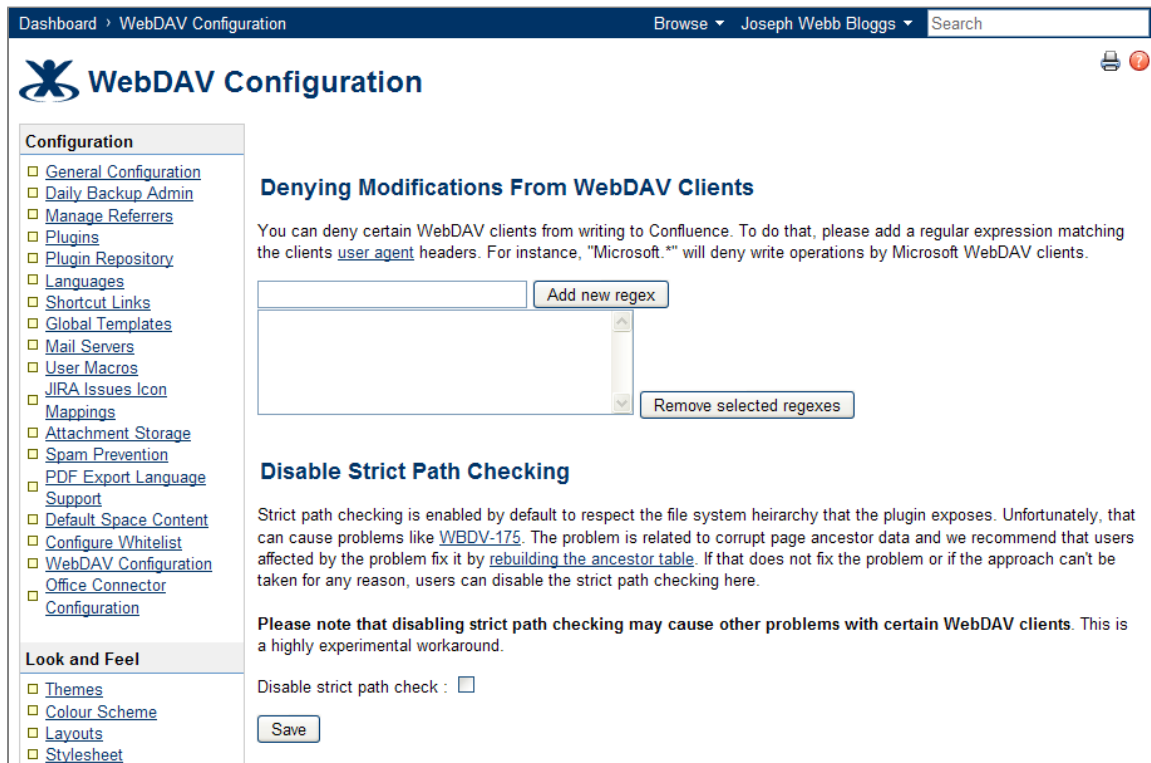
To restrict a WebDAV client's write access permissions to your Confluence installation,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Enter a regex that matches a specific component of the user agent header sent by the WebDAV client you want to restrict.
4. Click the '**Add new regex**' button. The regex is added to the list of restricted WebDAV clients.

   🛈 You can repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
5. Click the '**Save**' button to save the configuration changes.

To restore one or more restricted WebDAV client's write access permissions to your Confluence installation,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Select the regex(es) from the list that match(es) the user agent header sent by the restricted WebDAV client(s) you want to restore.
4. Click the '**Remove selected regexes**' button. The regexes you had selected are removed from the list of restricted WebDAV clients.
5. Click the '**Save**' button to save the configuration changes.

Screenshot: WebDAV configuration

Disabling Strict Path Checking

If you observe any idiosyncrasies with your WebDAV client, such as a folder that does exist on your Confluence site but is missing from the client, you can disable the WebDAV plugin's strict path checking option, which may minimise these problems.

To disable the WebDAV plugin's strict path checking option,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Clear the 'Disable strict path check' check box.

   🛈 You can re-enable this option at a later point in time by simply selecting this check box.
4. Click the '**Save**' button to save this configuration change.

Virtual Files and Folders

In the unlikely event that you observe any problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

> ⚠ By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, you must append the parameter `?hiddenOptionsEnabled=true` to the end of your URL and reload the page. For example:
>
> ```
> 1.<Confluence base URL>/admin/plugins/webdav/config.action?hiddenOptionsEnabled=true
> ```

Screenshot: The Hidden Virtual Files and Folders Option

To enable or disable access to virtual files and folders,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Amend your URL as described in the note above and reload the 'WebDav Configuration' page.
4. Select or clear the check box options in the 'Virtual Files and Folders' section as required.
5. Click the '**Save**' button to save the configuration changes.

RELATED TOPICS

📄 Attachment Storage Configuration (Confluence Docs 3.1)

📄 WebDAV Configuration (Confluence Docs 3.1)

📄 Important Directories and Files (Confluence Docs 3.1)

📄 Confluence WebDAV Plugin (Confluence Extensions)

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Other Settings

- Configuring Attachment Size
- Configuring Character Encoding
- Configuring HTTP Timeout Settings
- Configuring Indexing Language
- Configuring Jira Issues Icon mappings
- Configuring Number Formats
- Configuring Shortcut Links
- Configuring Time and Date Formats
- Number of Ancestors to Show in Breadcrumbs
- Thumbnail Settings

Configuring Attachment Size

Confluence gives you the option of limiting the maximum size of a single file attachment. Confluence administrators should keep in mind that the amount of disk space used by Confluence is directly proportional to the number and size of attachments put into the system.

To configure the maximum size allowed for an attachment

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' on the '**General Configuration**' screen.
3. Enter the maximum size next to '**Attachment Maximum Size**'. The default is 10 MB.
4. '**Save**' your changes.

To configure the maximum index-able size of attachments

By default, large attachment is defined as greater than 1 MB.
The threshold for attachments that won't get excerpts can be modified using the system property
`atlassian.indexing.contentbody.maxsize`, which takes a size in bytes.

Example

To specify 250 kb you would use the following JVM parameter:
`-Datlassian.indexing.contentbody.maxsize=256000`

Outcomes of Limiting Attachment Indexing Size

Limiting the size of attachment indexing has the following effects:

- Decreases the size of the index when large attachments are present.
- Decreases the memory used in indexing large attachments.
- Prevent excerpts of large attachments being displayed in search results.

For more details, please refer to the following JIRA issue.

RELATED TOPICS

Number of Ancestors to Show in Breadcrumbs

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

Configuring Time and Date Formats

Configuring Character Encoding

Thumbnail Settings

Configuring Attachment Size

Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Configuring Character Encoding

Confluence uses UTF-8 character encoding to deliver its pages.

 While it is possible to change the character encoding, unless you are certain of what you are doing, we recommend that you leave this as it is.

To change the character encoding,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Formatting and International Settings**' screen. For Confluence version earlier than 2.6.2, look for the '**Options and Settings**' screen.

3. Beside '**Encoding**', enter the new character encoding of your choice.

4. '**Save**' your changes.

RELATED LINKS

Joel Spolsky: The Absolute Minimum Every Software Developer Absolutely, Positively Must Know About Unicode and Character Sets (No Excuses!)

RELATED TOPICS

Number of Ancestors to Show in Breadcrumbs

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

Configuring Time and Date Formats

Configuring Character Encoding

Thumbnail Settings

Configuring Attachment Size

Recognised System Properties

Administrators Guide Home   Confluence Documentation Home

Configuring HTTP Timeout Settings

When macros such as the RSS Macro make HTTP requests to servers which are down, a long timeout value is used. You can set this timeout value through a system parameter to avoid this.

To configure the HTTP Timeout Settings,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' under the 'Configuration' heading in the left-hand panel.
3. Find the '**Connection Timeouts**' section in the lower portion of the screen.
4. Click '**Edit**' to adjust the settings.

Screenshot: HTTP Timeout Settings in Confluence

Connection Timeouts	Edit
External connections enabled :	true
Connection Timeout (milliseconds) :	10000
Socket Timeout (milliseconds) :	10000

HTTP Timeout Settings

The settings for HTTP Timeout in Confluence are as follows:

- **Adjust External connections enabled**: This setting allows system administrators to disable external connections so macros like the RSS Macro wont be allowed to make connections to an external server. It's provides protection against external servers providing insecure HTML, timing out or causing performance problems. The default setting is 'true'.

- **Connection Timeout (milliseconds)**: Sets the maximum time for a connection to be established. A value of zero means the timeout is not used. The default setting is ten seconds (10000).
- **Socket Timeout (milliseconds)**: Sets the default socket timeout (SO_TIMEOUT) in milliseconds, which is the maximum time Confluence will wait for data. A timeout value of zero is interpreted as an infinite timeout. The default setting is ten seconds (10000).

This feature is available in 2.2.8 and later versions of Confluence. Versions prior to 2.7 have a different method for adjusting these settings (link leads to legacy documentation).

Configuring Indexing Language

Changing the **Indexing Language** setting may improve the accuracy of Confluence search results if the majority of the content of your site is in some language other than English. Since Confluence v. 2.2.5, the content indexing support is provided in **German, Russian, Chinese, CJK, French, Brazilian, Czech and Greek besides English (default)**.

To configure a different indexing language,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' on the right hands side of the '**Formatting and International Settings**' section.

3. There is a drop-down list of '**Indexing Language**' currently supported in Confluence.

4. '**Save**' your changes.

RELATED TOPICS

Creating a Lowercase Page Title Index

Configuring Indexing Language

Rebuild the Content Indices from scratch

Content Index Administration

Working with Macros

Configuring Jira Issues Icon mappings

If you are using the {jiraissues} macro to retrieve information from a JIRA server, you will have to tell Confluence where to find the icons for any custom statuses or issue types you have configured in JIRA.

Confluence is configured by default with all JIRA's standard issue type and status icons. You will only need to change these settings if you have customised additional statuses or issue types for JIRA or have changed JIRA's default icons.

To configure custom icons,

1. Go to the '**Administration Console**' and click on '**Jira Issue Icon Mappings**' in the left panel.

2. For each icon you wish to configure, enter the *name* of the issue type or status into the **Jira entity** field, and the filename of its icon into the **filename** field.

    Ensure that the icon with that filename is located in the `/images/icons` directory of the JIRA server.

3. You may edit existing icon mappings by clicking on the `remove` link by an existing mapping, then re-adding it with a new icon filename.

RELATED TOPICS

Number of Ancestors to Show in Breadcrumbs

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

📄 Configuring Time and Date Formats

📄 Configuring Character Encoding

📄 Thumbnail Settings

📄 Configuring Attachment Size

📄 Recognised System Properties

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Configuring Number Formats

To change the number formats,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.

   - There are two number format settings:
     - Long Number Format
     - Decimal Number Format

3. Change the formats using the guidelines in this document.

4. '**Save**' your changes.

RELATED TOPICS

📄 Number of Ancestors to Show in Breadcrumbs

📄 Configuring Number Formats

📄 Configuring HTTP Timeout Settings

📄 Configuring Indexing Language

📄 Configuring Time and Date Formats

📄 Configuring Character Encoding

📄 Thumbnail Settings

📄 Configuring Attachment Size

📄 Recognised System Properties

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Configuring Shortcut Links

Shortcut links provide a quick way of linking to resources frequently referenced from Confluence. When you create a shortcut link, you are assigning a key to a URL so that when a user edits Confluence documents they can type the key instead of the complete URL.

Here is an example:

Most Google searches look like this: `http://www.google.com/search?q=`. If you create a shortcut for this search with the key 'google', every time a user needs to use `http://www.google.com/search?q=`**`searchterms`**, they can just type [**`searchterms`**@google] instead.

Here is a screenshot showing the shortcuts currently defined on http://confluence.atlassian.com:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		Remove
jira	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=	JIRA Issue %s	Remove
googlegroups	http://groups.google.com/groups?q=		Remove
google	http://www.google.com/search?q=		Remove
dictionary	http://www.dict.org/bin/Dict?Database=*&Form=Dict1&Strategy=*&Query=		Remove

Shortcut links are added and maintained by Confluence administrators from the **Administration Console**.

To create a shortcut link,

1. Go to the '**Administration Console**' and click '**Shortcut Links**' in the left panel.
2. Enter a '**Key**' for your shortcut. This is the shortcut name a user will use to reference the URL.
3. Enter the '**Expanded Value**'. This is the URL for the link. You can use '%s' in the URL to specify where the user's input is inserted. If there is no '%s' in the URL, the user's input will be put at the end.
4. *(Optional. Available in Confluence version 2.3 and later.)* Enter a '**Default Alias**'. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
5. Click '**Save**'.

Using Shortcut Links

Specify in the link what should go on the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to...	Type this	Resulting URL	Demonstration
a JIRA issue	[CONF-1000@JIRA]	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000	CONF-1000
a Google search	[Atlassian Confluence@Google]	http://www.google.com/search?q=Atlassian+Confluence	Atlassian Confluence@Google

Shortcut links can have titles just like any other link:

To link to...	Type this	Resulting URL	Demonstration
Internet Movie Database	[Fight Club|tt0137523@IMDB]	http://us.imdb.com/Title?tt0137523	Fight Club
Google Cache	[Cached Homepage|www.atlassian.com@cache]	http://www.google.com/search?q=cache:www.atlassian.com	Cached Homepage

Deleting Shortcut Links

Once you have created a shortcut link, it is listed under '**Shortcut Links**' in the '**Administration Console**'. Click '**Remove**' to delete the shortcut.

 RELATED TOPICS

Administrators Guide Home Confluence Documentation Home

Configuring Time and Date Formats

Confluence allows you to localise the formats used to display dates and times within the web interface. The settings use the syntax of Java's SimpleDateFormat class (described below).

To change the time and date formats,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.

   - There are three time and date format settings:
     - Time Format : displaying only the time of day (for example, when each news item is posted)
     - Date Time Format : displaying both the date and the time of day (for example, in historical versions of pages)
     - Date Format : displaying only the date (for example, the creation and most recent modification dates of pages)

3. Change the formats using the guidelines in this document.

4. '**Save**' your changes.

RELATED LINKS

- Java 1.4.2 SimpleDateFormat API

RELATED TOPICS

Number of Ancestors to Show in Breadcrumbs

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

Configuring Time and Date Formats

Configuring Character Encoding

Thumbnail Settings

Configuring Attachment Size

Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Number of Ancestors to Show in Breadcrumbs

Whenever there are three or more page links to be displayed in the breadcrumbs, Confluence will use an ellipsis like this '...' and display only the topmost and lowermost page links. Clicking on the ellipsis will display the page links in between.

 Note that the Dashboard and space homepage links are always displayed at the start of the breadcrumbs, and are not counted as ancestors for the purpose of this setting.

Screenshot: Breadcrumbs

Dashboard > Confluence 2.0 User Guide > ... > Viewing Page Information > Viewing History of a Page

You can configure how many immediate ancestors you want displayed in the breadcrumbs when you go to the page.

To configure the number of ancestors to show in the breadcrumbs,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.

2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.

3. Beside '**Number of Ancestors to show in Breadcrumbs**', enter a number. For example, if you enter 2, two immediate ancestors for the page will be displayed following the dots.

4. '**Save**' your changes.

Number of Ancestors to Show in Breadcrumbs

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Thumbnail Settings

The thumbnail settings allow you to define the height and width of images when they are displayed as thumbnails. This affects the images displayed by the Gallery macro and the Thumbnail macro.

To configure thumbnail settings,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.

2. Click '**Edit**' on the '**General Configuration**' screen.

3. Under the heading '**Attachment Settings**', enter a value in pixels for:
   - **Thumbnail maximum height** — The default setting is 200 pixels.
   - **Thumbnail maximum width** — The default setting is 200 pixels.

4. '**Save**' your changes.

RELATED TOPICS

📄 Gallery Macro

📄 Displaying a Thumbnail Image

📄 Uploading a Profile Picture

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Configuring System Properties

In general, you can configure system properties by providing an argument of `-Dprop=value` to a Java program when it starts up.

Because Confluence is a Java web application, the Java program is typically your application server. Therefore, you need to configure system properties in your application server's start-up script.

Below is a general example of how system properties are configured for any Java application. The system property arguments can appear anywhere in the argument list.

```
java ... -Dhttp.proxyHost=proxy.example.org -Dhttp.proxyPort=8080
```

The above example configures two system properties: `http.proxyHost` and `http.proxyPort` with values `proxy.example.org` and `8080` respectively.

Sometimes instructions will say simply 'set system property X', without saying what it should be set *to*. In this case, it is usually sufficient to simply use `-Dprop`, without a value. For example:

```
java ... -Datlassian.mail.disable
```

Application Server Examples

Please consult your application server documentation for how to provide system properties to the Java runtime or to a particular application server.

Below are some examples of how to start up different application servers with these system properties set:

AppServer	Startup Script	Variable to Edit	Example

Confluence Stand-alone .zip or .tar.gz Distribution	bin\setenv.bat (Windows)
	 bin/setenv.sh (Unix, Linux OS X)
		JAVA_OPTS	`set JAVA_OPTS=-Dhttp.proxyHost=proxy.example`
			 `-Dhttp.proxyPort=3128` (Windows)
			 `export JAVA_OPTS="-Dhttp.proxyHost=proxy.exampl`
			 `-Dhttp.proxyPort=3128"` (Unix)

Tomcat Windows Service	n/a	n/a	See below

Tomcat .war installation	bin\catalina.bat (Windows)
	 bin/catalina.sh (Unix)
		JAVA_OPTS	`set JAVA_OPTS=%JAVA_OPTS%`
			 `-Dhttp.proxyHost=proxy.example.com`
			 `-Dhttp.proxyPort=3128` (Windows)
			 `export JAVA_OPTS="-Dhttp.proxyHost=proxy.exampl`
			 `-Dhttp.proxyPort=3128"` (Unix)

Weblogic	$BEA_HOME/domain/<domain>/startWebLogic.sh	JAVA_OPTIONS	`export JAVA_OPTIONS="$JAVA_OPTIONS`
			 `-Dhttp.proxyHost=proxy.example.com`
			 `-Dhttp.proxyPort=3128"`

JBoss	bin/run.sh	JAVA_OPTS	`export JAVA_OPTS="-Dhttp.proxyHost=proxy.exampl` `-Dhttp.proxyPort=3128"`
WebSphere	n/a	Application Server > Process Definition > JVM Arguments	-Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128

Configuring System Properties in Tomcat as a Windows Service

Defining JAVA_OPTS in your `setenv.bat` file will not be sufficient to configure system properties if you are running Confluence with Tomcat as a Windows Service. Rather, there are two ways to configure system properties:

- through the Tomcat configuration application, normally found in the Windows system tray
- updating the service configuration directly in your Windows Registry.

As an example of the latter, for Tomcat 5 you will need to modify `HKEY_LOCAL_MACHINE -> SOFTWARE -> Apache Software Foundation -> Procrun 2.0 -> Tomcat5 -> Parameters -> Java -> Options` and add the JAVA_OPTS parameter there with the necessary system properties. JvmMx and JvmMs are listed separately from the additional options. See Editing the Windows Registry for details.

Displaying the System Properties

To see what Confluence is using, check Displaying System Properties.

RELATED TOPICS

Recognised System Properties

Recognised System Properties

Confluence has a small number of obscure configuration and debugging settings that can be enabled through Java system properties. System properties are usually set by passing the `-D` flag to the Java virtual machine in which Confluence is running. (Refer to the full instructions.)

Property	Since	Default Value	Module...	Effect

`atlassian.forceSchemaUpdate`	1.0	`true`	atlassian-config	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
`confluence.home`	1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the `confluence-init.properties` file, and use this property as the setting for the Confluence Home directory.
`confluence.devmode`	1.0	`true`	Confluence	Enables additional debugging options that may be of use to Confluence developers. Do not enable this flag on a production system.
`confluence.disable.mailpolling`	2.4	`false`	Confluence	If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail
`confluence.i18n.reloadbundles`	1.0	`true`	Confluence	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalised string is looked up. This can be useful when testing translations, but will make Confluence run *insanely slowly*.
`confluence.ignore.debug.logging`	1.0	`true`	Confluence	Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message.
`confluence.jmx.disabled`	3.0	`false`	Confluence	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in `WEB-INF/classes/jmxContext.xml`
`confluence.optimize.index.modulo`	2.2	`20`	Confluence	Number of index queue flushes before the index is optimised.
`confluence.plugins.bundled.disable`	2.9	`false`	Confluence	Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system.
`atlassian.mail.fetchdisabled`	1.0	`false`	atlassian-mail	Disables mail fetching services for IMAP and POP
`atlassian.mail.senddisabled`	1.0	`false`	atlassian-mail	Disables sending of mail
`atlassian.disable.caches`	2.4	`true`	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in devlopment if you are frequently changing static resources and don't want to continually flush your browser cache.
`org.osgi.framework.bootdelegation`	2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler. **,com.yourkit.**".
`confluence.diff.timeout`	3.1	1000	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.

RELATED TOPICS

Configuring System Properties

Configuring Logging

We recommend that you configure Confluence's logging to your own requirements. You can change the log settings in two ways:

- Configure logging in Confluence Administration – Your changes will be in effect only until you next restart Confluence.
- Edit the properties file – Your changes will take effect next time you start Confluence, and for all subsequent sessions.

Both methods are described below.

Terminology: In log4j, a 'logger' is a named entity. Logger names are case-sensitive and they follow a hierarchical naming standard. For example, the logger named `com.foo` is a parent of the logger named `com.foo.Bar`.

Configure logging in Confluence Administration

You can change some of Confluence's logging behaviour via the **Administration Console** while Confluence is running. Any changes made in this way will apply only to the currently-running Confluence lifetime. The changes are not written to the `log4j.properties` file and are therefore discarded when you next stop Confluence.

Not all logging behaviour can be changed via the Administration Console. For logging configuration not mentioned below, you will need to stop Confluence and then edit the logging properties file instead.

The '**Logging and Profiling**' screen shows a list of all currently defined loggers. On this screen you can:

- Turn page profiling on or off.
- Turn detailed SQL logging on or off.
- Add a new logger for a class/package name.
- Remove a logger for a class/package name.
- Set the logging level (INFO, WARN, FATAL, ERROR or DEBUG) for each class or package name.
- Reset all logging levels to a predefined profile.

Changing the logging configuration

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.

    ℹ️ You need to have System Administrator permissions in order to perform this function.
3. The '**Logging and Profiling**' screen appears, as shown below. Use the following guidelines to change the logging behaviour while Confluence is running:
    - '**Performance Profiling**' — See Page Request Profiling.
    - '**SQL Logging**' — Click the '**Enable SQL Logging**' button to log the details of SQL requests made to the database.

        ℹ️ If you need to enable logging of SQL parameter values, you will need to change the setting in the properties file. This option is not available via the Administration Console.
    - '**Log4j Logging**' — Click one of the profile buttons to reset all your loggers to the predefined profiles:
        - The '**Production**' profile is a fairly standard profile, recommended for normal production conditions.
        - The '**Diagnostic**' profile gives more information, useful for troubleshooting and debugging. It results in slower performance and fills the log files more quickly.
    - '**Add New Entry**' — Type a class or package name into the text box and click the '**Add Entry**' button. The new logger will appear in the list of '**Existing Levels**' in the lower part of the screen.
    - '**Existing Levels**' - These are the loggers currently in action for your Confluence instance.
        - You can change the logging level by selecting a value from the '**New Level**' dropdown list. Read the Apache documentation for a definition of each level.
        - Click the '**Remove**' link to stop logging for the selected class/package name.
4. Click the '**Save**' button to save any changes you have made in the '**Existing Levels**' section.

Screenshot: Changing Log Levels and Profiling

Performance Profiling

Profiling is currently OFF.

[Enable Profiling]

SQL Logging

[Enable SQL Logging]

Log4j Logging

Choose from one of the predefined logging options or configure logging below.

[Production]   [Diagnostic]

OR:

Customise specific logging settings

Add New Entry

Class/Package Name	New Level	
	INFO ▼	[Add entry]

Existing Levels

Class/Package Name	Current Level	New Level	
com.atlassian.confluence.cluster	INFO	INFO ▼	Remove
com.atlassian.confluence.cluster.safety	INFO	INFO ▼	Remove
com.atlassian.confluence.importexport.impl.PdfExporter	ERROR	ERROR ▼	Remove
com.atlassian.confluence.lifecycle	INFO	INFO ▼	Remove
com.atlassian.confluence.upgrade	INFO	INFO ▼	Remove
com.atlassian.core.util.FileUtils	ERROR	ERROR ▼	Remove
com.atlassian.upgrade	INFO	INFO ▼	Remove
net.sf.hibernate.cache.ReadWriteCache	ERROR	ERROR ▼	Remove
net.sf.hibernate.impl.SessionImpl	ERROR	ERROR ▼	Remove
net.sf.hibernate.type.CustomType	ERROR	ERROR ▼	Remove
net.sf.hibernate.util.JDBCExceptionReporter	ERROR	ERROR ▼	Remove
org.apache.fop	ERROR	ERROR ▼	Remove
root	WARN	WARN ▼	Remove

[Save]

Editing the Properties File

To configure the logging levels and other settings on a permanent basis, you need to stop Confluence and then change the settings in the `log4j.properties` file, described above.

The properties file contains a number of entries for different loggers that can be uncommented if you are interested in logging from particular components. Read more in the Apache log4j documentation.

See Working with Confluence Logs for some guidelines on specific configuration options you may find useful.

External Gadgets

The External Gadgets section allows you to register gadgets from other (external) web applications, such as JIRA 4.0+, iGoogle or Gmail, with Confluence so that they:

- Appear in the macro browser
- Can be added and used in Confluence pages or blog posts via a gadget macro

On this page:

- Obtaining the External Gadget's URL
- Registering an External Gadget for Use in Confluence
- Removing Access to an External Gadget in Confluence

Obtaining the External Gadget's URL

Before registering an external web application's gadget with Confluence, you will need to obtain that gadget's URL and copy it to your computer's clipboard.

If your web application is another Atlassian application such as Confluence 3.1+ or JIRA 4.0+, please refer to the appropriate documentation to obtain the gadget URL from your other Atlassian application:

- Obtaining a gadget URL from JIRA 4.0+
- Obtaining a gadget URL from another Confluence 3.1+ server

If your external gadget comes from a non-Atlassian web application or web site, please consult the relevant documentation for that application to obtain the gadget's URL.

Registering an External Gadget for Use in Confluence

To register an external web application's gadget for use in Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click 'External Gadgets' under 'Configuration' in the left panel. The 'External Gadgets' page is displayed.
3. In the 'Add a new Gadget' section, paste your gadget's URL into the '**Gadget Specification URL**' field.
4. Click the '**Add**' button. Your gadget will be shown in the list of registered gadgets below and it will also become available in the macro browser.

Do I need to establish an OAuth or Trusted Application relationship too?

In addition to registering an external gadget for use in Confluence, you may also need to establish an OAuth or Trusted Application relationship between the application that serves the gadget (the service provider) and Confluence (the consumer). OAuth and Trusted Application relationships are usually only required for gadgets that access user-restricted data from the external web application. Refer to Configuring OAuth for further information.

If an external web application provides anonymous access to all or some of its data and that is the only data you need to access, then establishing an OAuth or Trusted Applications relationship may be unnecessary.

Removing Access to an External Gadget in Confluence

To remove Confluence's access to an external web application's gadget,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click 'External Gadgets' under 'Configuration' in the left panel. The 'External Gadgets' page is displayed.
3. In the 'Added Gadgets' section, click '**Remove**' next to the URL of the external gadget whose access in Confluence is to be removed. The gadget will be removed from the 'Added Gadgets' list and will also be removed from the macro browser.

Screenshot: External Gadgets page

Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.

You can add gadgets from Atlassian applications such as Confluence, JIRA and others. You can also add gadgets from other websites such as iGoogle. Many public gadgets will work on a Confluence page. Some gadgets may rely on specific OpenSocial features that will not work properly in Confluence.

A gadget's URL looks something like this: http://example.com/my-gadget-location/my-gadget.xml

Add a new Gadget

Gadget Specification URL

[   ]

[Add]

Added Gadgets

Gadget Specification URL

https://jira.bamboo.atlassian.com/rest/gadgets/1.0/g/com.atlassian.bamboo.gadgets/charts/bambooCharts.xml	Remove
https://bamboo.extranet.atlassian.com/rest/gadgets/1.0/g/com.atlassian.bamboo.gadgets/charts/bambooCharts.xml	Remove
https://jira.bamboo.atlassian.com/rest/gadgets/1.0/g/com.atlassian.bamboo.gadgets/charts/planSummaryChart.xml	Remove
https://bamboo.extranet.atlassian.com/rest/gadgets/1.0/g/com.atlassian.bamboo.gadgets/charts/planSummaryChart.xml	Remove

Confluence and JIRA

- Add Confluence EAR-WAR to JIRA Standalone
- Integrating JIRA and Confluence
- Override properties in JIRA to Confluence Bridge
- Setting Up Trusted Communication between JIRA and Confluence

RELATED TOPICS

- Configuring Jira Issues Icon mappings
- JIRA Issues Macro
- JIRA Portlet Macro

Add Confluence EAR-WAR to JIRA Standalone

⚠ This guide is for experts only. If you run into any difficulties with this process, Atlassian technical support may provide limited assistance outside of helping users switch to running Confluence Standalone separately.

This document will assist you in adding Confluence to your existing **JIRA Standalone** instance.

Step 0 - Consider alternatives

Before embarking on this process, consider whether you could not rather run JIRA and Confluence in separate Tomcat instances running behind an Apache frontend server (see guides for Confluence and JIRA). There are some benefits to keeping them separate:

- Each app can be restarted without affecting the other.
- If one webapp hangs for any reason (eg. running out of memory), it doesn't affect the other.
- Any problems can be debugged more easily. Logs are separate and product-specific, rather than everything going to catalina.out.

Thread and heap dumps are smaller and more relevant.
- It reduces the likelihood of jar conflicts (eg. jars that must be installed in `common/lib` or `lib` for Confluence running off Apache Tomcat version 6 or above), particularly if you later want to install a third webapp not from Atlassian.

Offsetting this is the extra complexity of having to run Apache. The advantage of running two war files inside one application server is a moderate memory saving based on one jvm instead of two.

If you wish to proceed, please follow these instructions:

Step 1 - Download and extract WAR

1. Download the Confluence WAR file
2. Extract the downloaded zip file. It should extract to a folder called confluence-<version>. Inside this folder you'll find a folder called "confluence". Make a note of the absolute path to this directory (as you will need to use it later). *Note: Do not copy the confluence folder to the webapps folder inside tomcat - this may cause Confluence to be deployed more than once.*

Step 2 - Configure confluence-init.properties

1. Open `confluence/WEB-INF/classes/confluence-init.properties` in a text editor
2. Set the `confluence.home` property to a directory of your choosing. This is the directory that will contain all of Confluence's configuration, backup and attachment files.

Step 3 - Edit tomcat context descriptors

If you are deploying to **JIRA version 3.3 or higher**:

1. Create a file called *confluence.xml* in your JIRA standalone's *conf/Catalina/localhost* directory (if you have set up a different hostname for your JIRA tomcat instance, please specify that instead of *localhost*)
2. Open *confluence.xml* and add these lines:

```
1.<Context path="/confluence" docBase="c:/applications/confluence-2.1.3/confluence" debug="0"
reloadable="true">
2.          <Logger className="org.apache.catalina.logger.FileLogger" prefix=
"atlassian-confluence." suffix=".log" timestamp="true"/>
3.</Context>
```

3. For *docBase* specify the value you noted down earlier. This is the full path to the confluence folder in your confluence-<version> installation folder, **not** the confluence home folder. It should look like: `c:/<path to confluence installation>/confluence-<version>/confluence`.

Otherwise (**for older versions of JIRA**):

1. Open `conf/server.xml` in a text editor
2. Find the block that begins: <Context path="" docBase="../atlassian-jira" debug="0" reloadable="true"> and ends with </Context> block.
3. **After** the `</Context>`, append the following:

```
1.<Context path="/confluence" docBase="c:/applications/confluence-2.1.3/confluence" debug="0"
reloadable="true">
2.          <Logger className="org.apache.catalina.logger.FileLogger" prefix=
"atlassian-confluence." suffix=".log" timestamp="true"/>
3.</Context>
```

4. For *docBase* specify the value you noted down earlier. This is the full path to the confluence folder in your confluence-<version> installation folder, **not** the confluence home folder. It should look like: `c:/<path to confluence installation>/confluence-<version>/confluence`.
5. Remove the **commons-logging-1.0.4.jar** file from the confluence\WEB-INF\lib directory

Do not delete the existing Jira <Context> block. Insert the code above **after** the Jira <Context> block.

Step 4 - Modify your setenv.sh/bat (ONLY IF YOU ARE RUNNING A SUN MICROSYSTEMS JVM)

1. Open JIRA's bin/setenv.sh/bat (.sh on unix, .bat on windows) in a text editor.
2. Find the line that says: "... JAVA_OPTS=... " and add **-XX:MaxPermSize=128m** to its end.

The Java Virtual Machine sets aside a portion of memory as the "permanent space", for objects that it never expects to have to garbage-collect. Because JIRA and Confluence are both quite large applications, it is possible that this permanent space will be filled up. Increasing the application heap size will not help, because the permanent generation size is a separate setting.

> ⚠ This flag is only supported on JVM's created by Sun Microsystems. If you include this flag while running another vendors JVM (such as JRocket, or IBM's JVM) they will not start.

Step 5 - Restart the Server

1. Shut down, and then restart the standalone server
2. Confluence should now be accessible on the same server as your existing JIRA standalone, under the `confluence` directory. For example, if your JIRA is running at http://jira.example.org:8080/, Confluence will be running on http://jira.example.org:8080/confluence.

⚠ When setting up the Confluence database **do not** reuse the JIRA database. Create a new database for Confluence.

Troubleshooting

When I try to send a test mail from Confluence, I get `javax.mail.NoSuchProviderException: smtp`

In some circumstances, Confluence will be unable to send email after being deployed in the same application-server as JIRA. If, when you try to send a test mail from the administration tool, you get the error: "javax.mail.NoSuchProviderException: smtp", please follow these instructions to fix it.

I have installed JIRA and Confluence in some other application server than Tomcat

These instructions only apply to the standalone Tomcat version of JIRA. Other application servers have not been tested in this configuration, and users have specifically reported difficulties deploying the two applications together in Orion Server and JBoss. We hope to resolve these issues soon.

Confluence is slow, and dumps enormous amounts of information to its logfiles

If you are seeing a large amount of DEBUG logging output, then ensure that you have removed the **commons-logging-1.0.4.jar** file from the confluence\WEB-INF\lib directory

Integrating JIRA and Confluence

ⓘ Take a look at the technical guide to the process of adding your Confluence EAR-WAR to JIRA Standalone
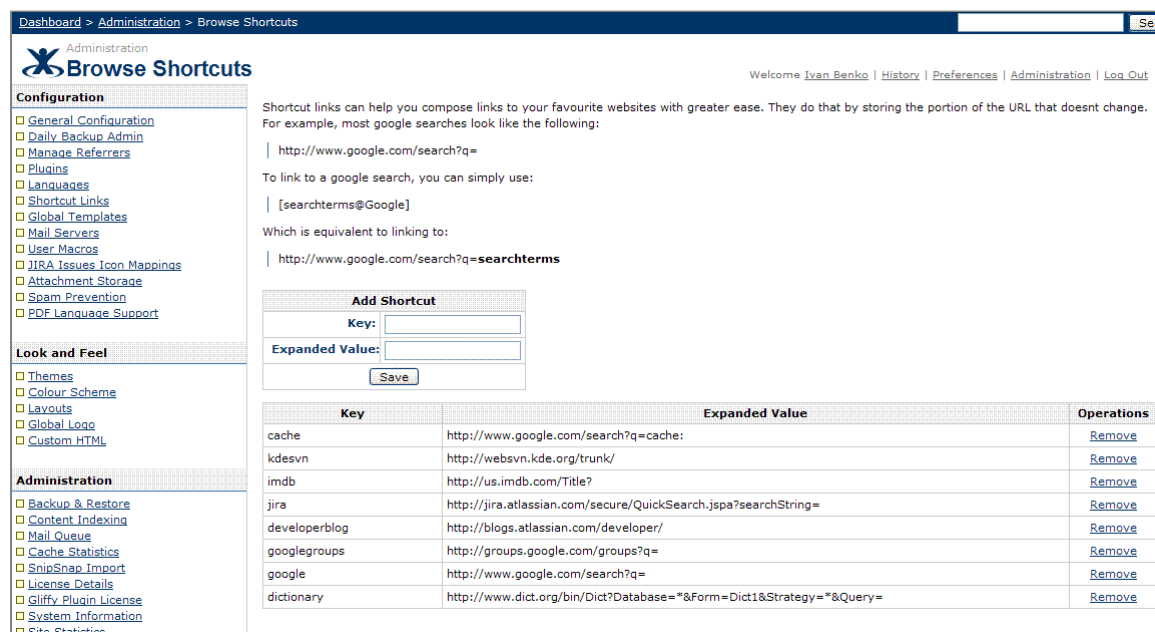
JIRA and Confluence were designed to complement each other. We've all seen projects where people try to store all their knowledge in the issue tracker, and we've seen projects where people have suffered trying to track issues in a knowledge management tool. We say: collect your team's thoughts, plans and knowledge in Confluence, track your issues in JIRA, and let the two applications work together to help you get your job done.

Here are four ways you can get JIRA and Confluence working together: use Confluence shortcuts to make easy links to JIRA issues, use trackback for two-way linking between Confluence and JIRA, use macros to include JIRA reports in Confluence pages, and integrate your JIRA and Confluence user management.

Combine Confluence Shortcuts and JIRA Quick Search

The simplest ideas can often be the most useful. In our Confluence site's global configuration - `Administration > Shorcut Links`, we have the following shortcut defined:

```
JIRA: http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=
```


This way, it's simple to create links using Confluence's shortcut notation. Link directly to JIRA issues: CONF-1000, or use JIRA's intuitive quick-search functionality to create links to particular groups of issue: CONF open improvements will link to a list of all open issues in the

Confluence project of type "Improvement" (try it and see!)

Use Trackback for easy two-way linking

Activate Trackback in JIRA and Confluence, and if someone makes a link from one application to the other, the link will automatically lead both ways: create a link from a JIRA issue to an example in a Confluence page, and the Confluence page will automatically know to link back to the JIRA issue, and vice versa. This is the perfect way to keep discussion connected to an issue.

- Document your user stories or use-cases in Confluence, and see at a glance which issues affect each use-case.
- If a JIRA issue requires more discussion or thought than can be conveniently held in comments, link them to a Confluence page.

(Note: as of Confluence 1.0 and JIRA 2.6, there is no mechanism for trackback to log in to JIRA or Confluence, so the use of trackback is limited to pages that are visible to anonymous visitors. In a protected Intranet environment, you may wish to open up Anonymous access to JIRA and Confluence to allow trackback to take place. Future revisions of the applications will give you the opportunity to allow Confluence to "log in" to JIRA and vice versa, avoiding this limitation)

Use the {jiraissues} and {jiraportlet} macros to embed JIRA reports and portlets into your Confluence site

Any JIRA search-result can be embedded in a Confluence page using the {jiraissues} macro with your choice of included fields and field ordering, and any JIRA dashboard portlet can be embedded in a Confluence page using the {jiraportlet} macro.

This way you can incorporate information from JIRA into the normal flow of your knowledge management. Combined with other macros like {junitreport}, {rss} and {html-include} and the FatCow suite, you can create dashboards in Confluence consolidating information from across your project, with Confluence and JIRA at the centre.

For **Confluence 2.7.0 and later**, an administrator can configure JIRA (3.12.0 or later) and Confluence to communicate in a trusted way, so that Confluence can request information from JIRA on behalf of the currently logged-in user. JIRA will not ask the user to log in again or to supply a password.

Trusted communication is used when embedding information from one application (e.g. a list of JIRA issues) into another application (e.g. a Confluence page).

Read more about trusted communication.

Link to Confluence pages from JIRA issues

While it is possible to simply paste links to Confluence pages into text fields of an issue (e.g. descriptions), the JIRA Linker Plugin provides a custom field that helps you find the correct page.

Integrate JIRA and Confluence user-management

To save you having to enter users into both JIRA and Confluence, you may benefit from using Atlassian Crowd as the user-repository for both applications.

Alternatively you can configure Confluence to use JIRA's user database (this requires that you are using JIRA with an external database; it will not work if you are using JIRA with an embedded HSQL database).

Some useful extensions

1. JIRA Confluence portlet - Display a Confluence page on the JIRA dashboard.
2. Atlassian Activity Stream Plugin - Activity Stream collects information from JIRA, Confluence, FishEye and Crucible.
3. AppLinks Plugin - Allows you to link projects, spaces and repositories between JIRA, Confluence, FishEye, Crucible and SVN applications without the need for long URLs.

And much more coming...

When you buy a license for JIRA or Confluence, you are automatically entitled to a year of updates. We listen to our customers needs, and having our products complement and work well with each other is very important to us. So if there is any way you think Confluence and JIRA could be made to work better, suggest it in our discussion space, and it may very well end up in a future version.

You might also like to take a look at our beyond JIRA page or watch the short video overview on some of these points in .mov format.

Override properties in JIRA to Confluence Bridge

Overriding properties used in the JIRA and Confluence Bridge

If, for some reason, you need to override the name of a column or a table used in Confluence's bridge to JIRA, you may do so in osuser.xml (see below).

This is most likely something you would consider doing if columns names were failing because your database is case sensitive.

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.<property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
03.                    <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
04.           </provider>
05.           <provider class="bucket.user.providers.CachingAccessProvider">
06.                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
07.                    <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
08.           </provider>
09.           <provider class="bucket.user.providers.CachingProfileProvider">
10.                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
11.                    <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
12.                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
13.           </provider>
```

Simply add the **<property name="chain.PROPERTY_NAME_HERE">NEW_VALUE</property>** element, to override a property (see below) with a new value.

Name Value Pairs for JiraJdbcAccessProvider, JiraJdbcProfileProvider and JiraJdbcCredentialsProvider

Property	Default value
userTable	userbase
userName	userName
userPassword	password_hash
groupTable	groupbase
groupName	groupname
membershipTable	membershipbase
membershipUserName	user_name
membershipGroupName	group_name
userId	id
membershipId	userId

Setting Up Trusted Communication between JIRA and Confluence

An administrator can configure JIRA and Confluence to communicate in a trusted way, so that Confluence can request information from JIRA on behalf of the currently logged-in user. JIRA will not ask the user to log in again or to supply a password.

When JIRA is configured to trust Confluence in this way, we call Confluence the '**trusted application**' and JIRA the '**trusting application**'.

Trusted communication is used when embedding information from one application (e.g. a list of JIRA issues) into another application (e.g. a Confluence page). Currently only JIRA can be configured to trust Confluence, and only the following two macros have been enhanced to use trusted communication:

- JIRA Issues macro
- JIRA Portlet macro

Further implementations will follow, especially as we roll out the tight integration required between Atlassian products for JIRA Studio.

> ⛔ **Potential security risk**
> Do not configure a trusted application unless you trust all code in that application to behave itself at all times. Trusted communication uses public/private key cryptography to establish the identity of the trusted server, so you must also be sure that the trusted application will maintain the security of its private key. Read the details of the security risks below.

On this page:

- Prerequisites
- Why do we need Trusted Communication?
- Overview
- Configuring JIRA to Trust Confluence

- Configuring the Macro Plugin in Confluence
- Adding the Macro to a Confluence Page
- Viewing the Confluence Page
- Security Risks
- Troubleshooting
- Technical Overview of the Trusted Applications Authentication (TAA) Protocol

Prerequisites

- JIRA 3.12.0 or later.
- Confluence 2.7.0 or later.
- In order to authenticate successfully against JIRA, the Confluence user must also be registered as a JIRA user with the same username.

> ⚠ **Common user base recommended**
> It is **highly recommended** that your JIRA and Confluence instances share a common user base, rather than two separate user bases with duplicated usernames. You will receive an error if Confluence passes JIRA a username which JIRA cannot recognise. Also, with separate user bases you run the risk that the same username may be used by two different people. The trusted application does not supply the user's password, so the trusting application will assume the username belongs to the user registered in the trusting application's own user base.

✅ **Tip:** Try Atlassian Crowd for a tidy user management solution.

Why do we need Trusted Communication?

The JIRA Issues and the JIRA Portlet macros allow you to embed a list of JIRA issues into a Confluence page. Prior to Confluence 2.7, if you wanted to display JIRA issues that had restricted viewing, then you needed to store the JIRA user's credentials (username and password) in the macro code directly on the Confluence page. This was not very secure.

The reasons we require the user credentials are:

- Your JIRA instance might not be public, and you might not want to allow anonymous access to your issues.
- You might have security restrictions on some of your issues. So you don't want to allow someone to leak data from your JIRA project by using the JIRA Issues Macro on a Confluence page.

Overview

Here is a summary of the integration points in a trusted communications relationship. Each of the following points is described in more detail in the sections below.

- A JIRA System Administrator configures JIRA to trust Confluence.
- A Confluence System Administrator configures the macro plugin to use (or not use) trusted communication.
- A Confluence user adds one of the macros to a Confluence page.
- A Confluence user or anonymous user views the Confluence page.

Configuring JIRA to Trust Confluence

Trust only has to be established once between the two applications. Once trust has been established, it is entirely transparent to the Confluence users.

Using the JIRA Administration Console, the JIRA System Administrator defines Confluence as a trusted application by specifying the Confluence instance's URL and other information. Refer to the JIRA documentation for details.

Configuring the Macro Plugin in Confluence

By default, Confluence ships with trusted communication enabled for the following macros:

- JIRA Issues macro
- JIRA Portlet macro

A Confluence System Administrator can decide on the level of trusted communication used by the macros. The different levels are:

- Ignore trusted communications altogether. Trusted communication is turned off at the global level.
- Perform trusted communications whenever the macro is used on a Confluence page, but do not show certain warning messages.
- Perform trusted communications whenever the macro is used on a Confluence page, and show all warning messages. **This is the default configuration.**

To change the default trusted communication level for the JIRA Macros plugin,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel.
3. The '**Plugin Manager**' screen appears, showing a list of installed plugins. Scroll down and click the '**JIRA Macros**' link.
4. The '**JIRA Macros**' panel appears in the top middle of the screen, as shown below. Click '**Enable**' or '**Disable**' next to the following options:
    - '**JIRA application trust support**' – With this option enabled, Confluence will attempt trusted communication with JIRA whenever a user views a page containing the JIRA Issues or Portlet macro, provided criteria are met as described below. With this option disabled, Confluence will never attempt trusted communication with JIRA for these macros.

        ✅ Disable the above option if you do not intend to configure trusted communication between JIRA and Confluence.
    - '**JIRA application trust warnings**' – With this option enabled, Confluence will display all error and warning messages that may arise from a problem during trusted communication (assuming that trusted communication is enabled). With this option disabled, Confluence will suppress certain warnings. See troubleshooting below.

        ✅ Disable the above option if you have a large number of existing JIRA macros already on your Confluence instance, pointing at a diverse range of JIRA servers. Some of those JIRA servers may have a trusted communication link established (requiring the functionality to be enabled) while other JIRA servers may have no trusted communication link. In this case, you may want to turn off the warning messages so they do not appear on your Confluence pages where the JIRA macros point to non-trusting JIRA servers.

Screenshot: JIRA Macros panel in Plugin Manager

Adding the Macro to a Confluence Page

The Confluence user can add and edit the macros as described on the following pages:

- Using the JIRA Issues macro
- Using the JIRA Portlet macro

> ✅ **Remove the username and password from your macro markup code**
> Prior to Confluence 2.7, you needed to include a username and password in the macro markup code if you wanted to display JIRA issues which had restricted viewing. Once your administrator has set up trusted communication between Confluence and JIRA, you no longer need to include a username and password in the markup code for your JIRA macros.

The following options are available for determining the issues which will be retrieved from JIRA and displayed on the Confluence page:

What you want to do	Macro parameter	URL parameter	Comments
Display the JIRA issues which the logged-in user is authorised to see. And if the user is not logged in, display only issues which allow unrestricted viewing.			Do not specify any authentication parameters. In this case, the behaviour depends on the way your administrator has set up trusted communication between JIRA and Confluence. Here is a summary of the behaviour. If trusted communication is **enabled**, the authorisation will work seamlessly. When a logged-in user views your page, they will see only the JIRA issues they are allowed to see. And if they are not logged in, they will see only the issues which allow unrestricted viewing. If trusted communication is **disabled**, the Confluence page will show only the JIRA issues which allow unrestricted viewing.
Ensure that Confluence will display only the JIRA issues which allow unrestricted viewing.	`anonymous`		Regardless of who the user is (logged in or not), the Confluence page will show only anonymously-visible issues. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
Use a pre-determined username and password to access the JIRA issues.		`&os_username=MYNAME&os_password=MYPASSWORD`	**Not recommended**. Prior to Confluence 2.7, this was the only way of displaying issues with restricted viewing. For Confluence 2.7 and later, this method will still work. Confluence will not attempt to set up a trusted communication link with JIRA in this case.

Refer to the section below for details of what happens when a user views a Confluence page containing a JIRA macro.

Viewing the Confluence Page

When a user views a Confluence page which contains a JIRA Issues or JIRA Portlet macro, this is what happens:

- If the macro markup contains an explicit username and password in the URL parameter, Confluence will not request trusted communication with JIRA. Confluence will retrieve the JIRA issues which the specified username is authorised to see. This behaviour is the same as Confluence versions prior to 2.7.
- If the macro markup contains the `anonymous` parameter, Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If the user is anonymous (not logged in), Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If trusted communication is disabled via the Plugin Manager in Confluence, then Confluence will not request trusted communication with JIRA. So if there is no explicit username and password in the markup code, Confluence will retrieve only the JIRA issues which allow unrestricted viewing. This behaviour is the same as Confluence versions prior to 2.7.
- If trusted communication is enabled via the Plugin Manager in Confluence:
    - If the user is logged in, then Confluence attempts trusted communication with JIRA. Confluence sends the username to JIRA. JIRA returns a set of issues which that username is authorised to access, based on the JIRA user base and the JIRA groups and permissions. Confluence displays those issues on the page.
    - If JIRA or Confluence encounters a problem during the trusted communication process, an error message may appear on the Confluence page above the macro output – see troubleshooting below.

Security Risks

Please take the following considerations into account when setting up trusted communication:

- When you configure JIRA to trust an application, you are allowing the application to access JIRA in the name of a particular user. The trusted application passes JIRA the user's login name, but no other authentication information. JIRA does not request the user's password. By doing this, you are **bypassing JIRA's authentication mechanism**.
- Do not configure a trusted application unless you **trust all code in that application** to behave itself at all times.
- Trusted communication uses public/private key cryptography to establish the identity of the trusted server. The trusted application

needs to maintain the security of its private key. Confluence stores its private key in the database. **So you must be sure that the Confluence database is secure, and also any full backups of the database.**

- Be aware of the risks associated with using separate user bases, as explained above. **We strongly recommend a common user base between the trusted and trusting applications.**
- When configuring an application to trust another application, you should use a trusted network or SSL to **protect the sensitive information passed between the applications during the configuration procedure**. This will help to prevent man-in-the-middle attacks.

Troubleshooting

Below are the warning messages which may appear on your Confluence page, above the output of the JIRA Issues or JIRA Portlet macro.

Warning Message	Cause	Solution	Warning Message Can be Turned Off?
`javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: PKIX path building failed: sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid certification path to requested target`	JIRA is running over SSL	Add JIRA's SSL Certificate to the Java Keystore	**No**
`The JIRA server does not recognise your user name. Issues have been retrieved anonymously.`	The logged-in Confluence user is not registered in the JIRA user base.	Add the username to your JIRA user base. It is **highly recommended** that your JIRA and Confluence instances share a common user base.	**No**
`The JIRA server does not trust this Confluence instance for user authentication. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.`	Your JIRA instance has not been configured to trust your Confluence instance.	One of the following solutions: • Configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the `anonymous parameter` in all your JIRA Issues and JIRA Portlet macros.	Yes
`The JIRA server does not support trust requests. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.`	Your JIRA instance is not able to handle trusted communications (i.e. the JIRA version is earlier than 3.12.0).	One of the following solutions: • Download the latest version of JIRA and then configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the `anonymous parameter` in all your JIRA Issues and JIRA Portlet macros.	Yes

Consult Troubleshooting the JIRA Issues Macro and Trusted Applications for further troubleshooting.

Technical Overview of the Trusted Applications Authentication (TAA) Protocol

⊘ Read this section if you want a bit more information on the technical side of things.

Atlassian has developed its own protocol to set up trust between JIRA and Confluence. Below is a technical overview of the process.

Configuring JIRA to trust Confluence:

1. When the JIRA System Administrator provides the base URL of the Confluence instance, JIRA requests a trusted application authentication certificate from Confluence. The certificate contains Confluence's trusted application ID and public key (generated specifically for use with the TAA protocol).
2. JIRA validates the certificate and asks the System Administrator for a few extra details about the trust relationship, such as a name for the Confluence instance, timeout, allowed IP addresses and allowed request URLs.
3. JIRA stores all this information in the database.

Making a trusted request from Confluence to JIRA:

1. Confluence sends a web request to JIRA, appending additional headers to the request, including:
   - Timestamp (nonce) of the request + user name of the currently logged-in Confluence user, encrypted with a symmetric key (generated on the fly).
   - The symmetric key, encrypted with Confluence's private key.
   - Confluence's application ID (as displayed when trusted communication was established).
2. JIRA attempts to decode the encrypted headers, using the stored information about the relationship. It conducts the following checks to validate the request:
   - The trusted application ID refers to a valid trusted application.
   - The given username exists in the JIRA user base.
   - The agreed timeout has not expired.
   - The request originated from a trusted IP address.
   - The resource being requested matches those specified in the URL match list.
3. If any of these checks fails, a response is sent to Confluence indicating the reason for failure. Otherwise, JIRA will authenticate the specified user for the duration of the single request, and respond with the resources (i.e. the JIRA issues).

RELATED TOPICS

JIRA Issues Macro
JIRA Portlet Macro
Connect to LDAP, JIRA or Other Services Via SSL
Single Sign-on Integration with JIRA and Confluence
Troubleshooting the JIRA Issues Macro and Trusted Applications

Confluence Clustering Overview

It is possible to run Confluence in a clustered environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.

> ⚠ **Consider your options carefully before deciding on a clustered installation**
> While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information on the Cluster Checklist and then consult Atlassian support before making your final decision.

This page gives an overview and links to further pages with information on installing, configuring and administering a Confluence cluster.

Before Deciding to Run a Confluence Cluster

1. Read and consider the details on the Cluster Checklist.
2. Consider the difference between clustering for scalability and clustering for high availability (HA).
3. Contact Atlassian support for further information and advice.

Technical Overview

Read a technical overview of clustering in Confluence.

Server and Network Requirements

- Server hardware requirements
- Technical overview of Confluence clustering
- Diagram of recommended network topology

Installation and Upgrading

There are two methods of installing Confluence in a cluster, depending on whether you have existing data:

- Fresh installation
- Existing data

If you are upgrading an existing Confluence cluster to a new version of Confluence, refer to the cluster upgrade guide.

Configuration and Administration

- Cluster Administration page in the **Administration Console**
- Changing datasources in clusters

Troubleshooting

- Cluster troubleshooting

RELATED TOPICS

Operating Large or Mission-Critical Confluence Installations
Performance Tuning
Requesting Performance Support
Administrators Guide
Configuration Guide

Technical Overview of Clustering in Confluence

> **ⓘ** **Overview of clustering documentation**
> Refer to the overview of Confluence clustering in the *Administrators' Guide*.

Introduction

From version 2.3, Confluence has had the ability to configure and run multiple copies of itself in a cluster, so that clients can connect to any copy and see the same information. In effect, a Confluence cluster behaves as a single, powerful Confluence installation. While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change from earlier versions (or non-clustered installations) and consequently, requires extra planning for deployment and upgrades.

This document will give a technical overview of clustering in Confluence, primarily for those users and developers who will be installing and configuring Confluence in a cluster. A separate overview is available for Confluence plugin developers.

Cluster topology

A simple description of the cluster topology for Confluence would be **multiple applications, shared data source**. A cluster of Confluence consists of:

- multiple homogeneous installations of Confluence (called *nodes* below)
    - a Confluence home directory for each installation.
- a distributed Oracle Coherence cache (formerly known as Tangosol Coherence), which all nodes use via a multicast group - see networking summary below
- a single database, which all nodes connect to

The user is responsible for configuring an appropriate HTTP load balancer in front of the clustered installations. Typically this means using mod_jk or another application server load-balancing technology. The load balancer must be configured to support **session affinity**.

Communication between clustered nodes is minimised by using a distributed cache which propagates updates to all other nodes automatically. Where necessary, Coherence provides a locking mechanism for synchronising jobs and a RMI interface for more complex communication.

LAN Clustering Only

Atlassian **only supports clustering over a local area network**. While it is theoretically possible to configure Confluence to cluster across a WAN, the latency involved is likely to kill performance of the cluster. We can't stop you trying, of course, but you're going to have to work out how to configure Coherence yourself, and we're not going to support the resulting mess.
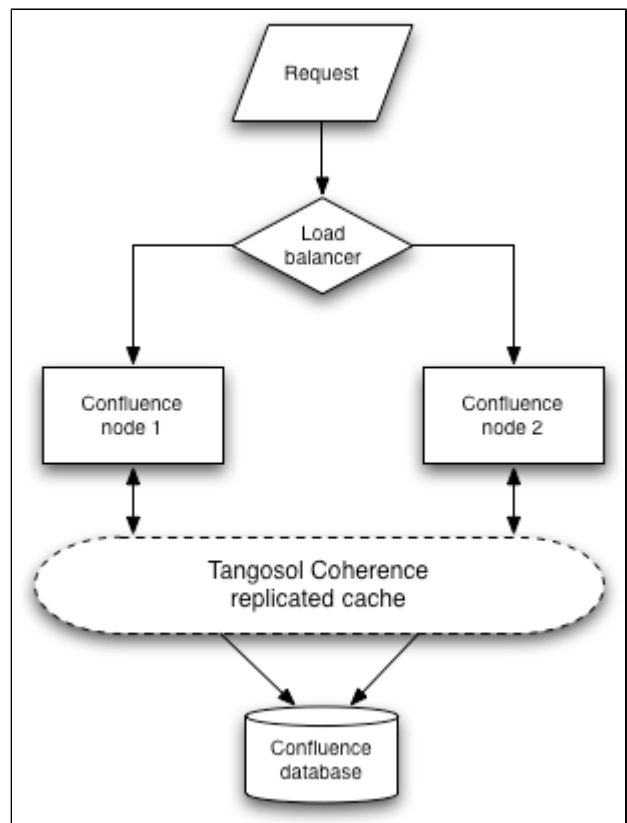
Confluence cluster topology (simplified)

Homogeneous Confluence installations

All the Confluence installations must be running exactly the same application, down to the lowest level. Items that must be the same include:

- Confluence version
- Application server version
- JDK version
- Libraries and plugins in the Confluence classpath, WEB-INF/lib
- Libraries in the application server classpath

The installation section has more information how to ensure homogeneous node installations.

Creating a Confluence cluster

When installing Confluence in a clustered setup, you will be responsible for configuring your web server and load balancer to distribute traffic between each node. No additional software is required as Coherence is bundled with Confluence.

Here is an overview of the process:

1. Obtain a clustered licence key from Atlassian for each node
2. Upgrade a single node to the clustered licence
3. Start the cluster from that node's administration menu, specifying a name and optionally a preferred network interface
4. Restart the single node and test it
5. Copy the Confluence application and Confluence home directory to the second node
6. Bring up the second node and it will automatically join the cluster.

Copying the Confluence application and home directory helps ensure that the installations are homogeneous.

An alternative to this method is to copy the Confluence web application, but not the Confluence home directory. In this case, the installation wizard will require your cluster name to connect to the other nodes, and it will automatically configure itself. You will need to rebuild the index manually after this installation, however.

There is now full documentation for a Confluence Cluster Installation.

Upgrade process

Another consequence of the homogeneous requirement is that upgrades must be done by following a strict process.

1. All cluster nodes are brought down
2. Upgrade a single node to the latest Confluence version
3. Start the single node so it can upgrade the database
4. Upgrade subsequent nodes and start them one-by-one.

This is the only safe method of upgrading a Confluence cluster.

Single database

The Confluence database in a cluster is shared by all nodes. This means that the database must be able to scale to service *all* the Confluence nodes, which will probably mean implementing some kind of database cluster and JDBC-level load balancing. We can not offer support with scaling or tuning your database, you will need to talk to your DBA or database vendor.

For obvious reasons, you must have an external database to run Massive - you can not cluster Confluence when using the embedded HSQL database.

The most important requirement for the cluster database is that it have sufficient connections available to support the expected number of application nodes. For example, if each Confluence instance has a connection pool of 20 connections and you expect to run a cluster with four nodes, your database server must allow at least 80 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

In a cluster, attachments must be stored in the database. Configuring a cluster in an existing installation will automatically migrate your attachments to the database. Non-clustered installations still have the option of using the Confluence home directory for storing attachments.

While attachments are stored in the database, they are temporarily written to the cluster node's local filesystem, designated `<confluence-home>/temp` folder, when being streamed to users (so Confluence doesn't have to hold open database connections unnecessarily). For this reason, Confluence will still need enough temporary disk space to hold any attachments currently in transit.

Distributed cache

In a normal configuration, Confluence uses many caches to reduce the number of database queries required for common operations. Viewing a page might require dozens of permissions checks, and it would be very slow if Confluence queried the database for this information with every page view. However, caches must be carefully maintained so they are consistent with the application data. If the page permissions change, the old invalid data needs to be removed from the cache so it can be replaced with a fresh correct copy.

To preserve consistent caches across a cluster, Confluence uses a distributed cache called Oracle Coherence, which manages replicating cache updates transparently across all nodes. The network requirements of the distributed cache are quite simple, but must be preserved if the cluster is to work properly.

To discover other nodes in the cluster, Confluence broadcasts a join request on a multicast network address. Confluence must be able to

open a UDP port on this multicast address, or it will not be able to find the other cluster nodes.

Once the nodes are discovered, each responds with a unicast (normal) IP address and port where it can be contacted for cache updates. Confluence must be able to open a UDP port for regular communication with the other nodes.

Because the Coherence network requirements are different to those required by the Confluence database connection, the situation can arise where Confluence can use the database but not talk to the other nodes in the cluster via Coherence. When Confluence detects this, it will shut itself down in a cluster panic.

For more details on the network configuration of the distributed cache, see the networking summary

Home directory

Confluence's home directory has a much-reduced role in a cluster. Because the application data must be shared between all nodes for consistency, the only information stored in the Confluence home directory is either node-specific, or needed to start Confluence. This includes information related to:

- database connection
- license
- cluster connection

The only application data stored in the Confluence home directory is the **Lucene search index**. Confluence synchronises this data itself by keeping track of indexing tasks in the database.

This is also why we recommend copying the Confluence home directory from the first node when setting up subsequent nodes. If you did not copy the Confluence home directory, you would need to rebuild the search index from scratch on the subsequent nodes after installation.

Event handling

Broadcasting events to all nodes in a cluster is supported in Confluence, but not recommended. The cluster topology uses a shared data store so that application state does not need to be synchronised by events.

The event broadcasting is done only for certain events, like installing a plugin. When a plugin is installed in one node, Confluence puts the plugin data in the database, and notifies the other nodes that they need to load the plugin into memory.

Indexing

Confluence maintains a copy of its Lucene search index on each node of the cluster. This index is used for many things beside full-text searches, including RSS feeds and lists of recently updated content. Indexing in a cluster works like this:

1. Node 1 gets a request to save some page update
2. After saving the page in the database, Node 1 adds a "page-updated" index entry to the queue, which is in the database
3. Periodically, each node picks up the "latest entries" from the queue, where what is latest is determined from a timestamp on a file in the Confluence home directory which indicates when the queue was last inspected. This process is called "flushing the index queue".
4. Each node independently updates its local Lucene index. The "page-updated" index entry is internally changed into a delete-document task and an add-document task to apply the changes to Lucene.
5. Each node updates the timestamp on its index-queue-timestamp file to reflect the most recent processing or "flushing" of the index queue.

Because of step #3, if the timing of the nodes is not synchronised or changes sporadically (due to a virtualisation environment, typically), index changes will not be correctly synchronised in the cluster. This is the most common cause of index sync problems in clusters.

If a node is disconnected from the cluster for a short amount of time (less than three hours), it will be able to bring its copy of the index up-to-date when it rejoins the cluster. If a node is down for a long amount of time and its lucene index has become stale as a result, you may want to avoid the expensive operation of rebuilding the index. To do that, you must copy a "live" version of the Lucene index from an active node. Simply replace the contents of the `Confluence Home]/index` directory with those from an active node before bringing the stale node back up.

Job synchronisation

For tasks such as sending the daily report emails, it is important that only one node in the cluster does this. Otherwise you would get multiple emails from Confluence every day.

Confluence uses locks in the Coherence distributed cache to ensure only one node can be running certain jobs at a time. This ensures email notifications will only be sent once.

Activity tracking

Activity tracking does not work in a cluster, and will be disabled for clustered deployments. We're working on making the activity tracker clusterable in a future release. You can follow this issue. You can try some other options for tracking usage.

Cluster panic

In some situations, there can be a network issue or firewall that prevents the distributed cache from communicating but still allows Confluence to update the database. This is a dangerous situation because when the caches on the detached nodes become inconsistent, users on different nodes will see different information and updates can be lost.

Confluence can detect this problem by checking a database value against a cached value, and if they differ, all the clustered nodes will be shut down with a 'Cluster panic' message. This is considered a fatal error because the consequences can cause damage to your data. For those administrators that like to live on the edge, there is a system property to prevent cluster panic and allow data corruption. For more information, see Cluster safety mechanism.

If a cluster panic does occur, you need to ensure proper network connectivity between the clustered nodes. Most likely multicast traffic is being blocked or not routed correctly. See the networking summary below.

Summary of network requirements

In addition to normal connectivity with its database, all clustered Confluence instances require access to a multicast group and the ability to open a UDP unicast port.

By default, the multicast address is automatically generated from the cluster name you provide when starting the cluster and the multicast port is fixed. During cluster setup, Confluence will prompt for the unicast IP address to use if the server has multiple network interfaces, and by default the unicast port is fixed. The cluster multicast group will be joined on the same network interface as the bound unicast IP address.

For any settings which are not configurable through the Confluence web interface, they can be configured via an XML file in the Confluence home directory for more exotic networking requirements.

Scaling Confluence On A Single Server

Since the maximum addressable memory on a 32 bit JVM is 4GB, some large servers may scale Java applications by running JVM instances concurrently. This would be implemented as separate, clustered Confluence nodes running on a single server and communicating internally. Because each JVM replicates the cache entirely, it may be useful to test a single, massive instance running a 64 bit JVM as an alternative. This configuration may result in superior performance than an internal cluster.

Geographically Distributed Clusters

Collocating nodes is strongly recommended as high latency will almost certainly degrade performance due to the overhead of cache replication. Cluster nodes will provide the best performance if servers are physically adjacent. However, as long as all nodes share a LAN, users may wish to test alternative configurations to see how performance is affected.

RELATED TOPICS

Server Hardware Requirements Guide
Overview of Confluence Clusters

Cluster safety mechanism

Introduction

A mechanism was added in Confluence 2.3 and above to ensure database consistency when running multiple cluster nodes against the same database. This is called the *cluster safety mechanism*, and is designed to ensure that your wiki cannot become inconsistent because updates by one user are not visible to another. A failure of this mechanism is a fatal error in Confluence and is called *cluster panic*.

Because the cluster safety mechanism helps prevents data inconsistency whenever any two copies of Confluence running against the same database, it is enabled in *all* instances of Confluence, not just clusters.

How cluster safety works

A scheduled task, ClusterSafetyJob, runs every 30 seconds in Confluence. In a cluster, this job is run only on one of the nodes. The scheduled task operates on a *safety number* – a randomly generated number that is stored both in the database and in the distributed cache used across a cluster. It does the following:

1. **Generate** a new random number
2. **Compare the existing safety numbers**, if there is already a safety number in both the database and the cache.
3. **If the numbers differ, publish a ClusterPanicEvent**. Currently in Confluence, this causes the following to happen:
    - disable all access to the application
    - disable all scheduled tasks
    - update the database safety number to a new value, which will cause all nodes accessing the database to fail.
4. If the numbers are the same or aren't set yet, **update the safety numbers**:
    - set the safety number in the database to the new random number
    - set the safety number in the cache to the new random number.

How to fix it

Cluster Panic

Usually presents itself with the following error message:

```
FATAL [DefaultQuartzScheduler_Worker-4] [confluence.cluster.safety.ClusterPanicListener]
handleEvent Fatal error in Confluence cluster:
Database is being updated by an instance which is not part of the current cluster. You should
check network connections between cluster nodes, especially multicast traffic.
```

In almost all cases, cluster panic events are caused by two or more instances of Confluence (in separate clusters) updating the same database. Such events are typically caused by one of the following issues:

JVM paused (e.g. while swapping memory) can break communication between two nodes

Always watch the swapping activity of your server and avoid swapping due to lack of RAM. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection (GC) considerably and affect Confluence's performance.

- In clustered installations, swapping can lead to cluster panic. This is because swapping causes the JVM to pause during garbage collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Two instances of Confluence have been started in your application server

This is one of the most commonly encountered issues. The strangest case of this that we have seen so far involved a cloned image of a PC running Confluence that was later used in a remote office in a different city. The people using Confluence on the cloned instance were not aware that the original Confluence instance was also running and that both these Confluence instances were using the same production database server.

- Solution: Check your application server's configuration to make sure that multiple copies of the application server are not running concurrently. Database transaction logs can help identify the location of other application servers, if client IP addresses are recorded along with each transaction.

Two copies of your application server are running.

Sometimes starting an application server twice will result in two processes running, even though only one can be accessed over the network.

- Solution: Check a list of running processes (for example, with the 'ps' command in Posix-based operating systems like Linux, Unix and Mac OS X) and make sure your application server is only running once.

Networking failure between nodes in the cluster

- Solution: Check that multi-cast traffic is being transmitted successfully, and that the network between your nodes is low-latency (<100 ms).

Database server stops responding

If Coherence fails to retrieve the SafetyNumber from the database, the comparison will fail. If it fails to update it, the next comparison will fail, 30 seconds later.
Many things can cause this, including a scheduled shutdown for backups, network failure, a filled-up transaction-log partition and a changed password on the account used by Confluence to connect to the database.

- Solution: resolve the problem with the database (or network), then restart Confluence

> 🚫 In all cases, when starting Confluence after a cluster panic, you must ensure all cluster nodes have been shut down completely. If necessary, use commands like `ps` and `kill` to get a list of Java processes and terminate them manually.

> ⚠ Please visit this document for troubleshooting advice if you encounter any of the above situations.

Technical details

The cluster safety number in the database is stored in the CLUSTERSAFETY table. This table has just one row: the current safety number.

Changing Datasources Manually in a Cluster

> 🚫 The recommended way of changing database connections is to shut down the whole cluster, install Confluence into new and empty directories and use the Setup Wizard to configure all new database connection settings.

However, if you wish to manually change your settings, you may proceed as described below.

⚠ It is **strongly recommended** that you test all of the following in a staging or test instance of Confluence before performing these steps in

your production environment.

Step 1: Prepare

- Locate the `confluence-cfg.xml` file in the Confluence home directory.
- Make a backup copy of that file.
- Prepare the necessary changes to that file.

Step 2: Shut Down Confluence

You need to shut down all the nodes in the cluster, not just one.

Step 3: Apply your Changes

Apply your configuration changes to the required node.

Step 4: Restart the Changed Node

It is crucial that you bring up the node on which you applied the changes **first**. Otherwise you will get an error message, and have to shut down all instances again.

Step 5: Restart all Other Nodes

Done.

RELATED PAGES

Overview of Confluence Clusters

Disabling a Cluster

Recommended Approach

Run the Migrating Confluence Between Servers procedure for a new stand-alone configuration with the data imported.

Possible Hack

Another option to run standalone mode (disable the clustered cache) with a cluster license is to set cluster property to false in the confluence.cfg.xml and restart confluence.

For example:
<property name="confluence.cluster">false</property>

Atlassian has tested this and it works, but it is not fully tested. As always, backup your database and confluence home directory before making any modification (as this is not the recommended approach).

To check to see if clustering has been disabled, look in the logs after the xml during startup.

In your catalina.out, you'll have:

```
1.INFO [main] [confluence.cluster.tangosol.TangosolClusterManager] startCluster Bringing up cluster service
```

This line won't exist if you start it up with the configuration above.

Cluster Troubleshooting

Clustering is vastly improved in Confluence 3.0. See the Confluence 3 Performance improvements page for more information. If you're having cluster performance problems in a version prior to Confluence 3, we suggest an upgrade!

Overview of clustering documentation
Refer to the overview of Confluence clustering.

On this page:

- Symptoms
- Confluence cluster debugging tools

Symptoms

Below is a list of potential problems with a Confluence cluster, and their likely solutions. The solutions are listed below.

Problem	Likely solutions
`Database is being updated by an instance which is not part of the current cluster` errors at startup	Add multicast route, Check firewall
`Cannot assign requested address` on startup, featuring an IPv6 address	Prefer IPv4
Error in log: `The interface is not suitable for multicast communication`	Change multicast interface, Add multicast route
Multicast being sent, but not received (detectable with Multicast Test)	Check firewall, Check intermediate routers, Increase multicast TTL
Any issue not covered here	Contact support

Confluence cluster debugging tools

There is an umbrella issue opened for all cluster debugging tools here

It includes the tools listed below.

Multicast

- Which multicast address?

The multicast address and port used by Confluence can be found on the Cluster Administration page, or in `confluence.cfg.xml` in the Confluence home directory.

- Multicast address generation.

Confluence uses a hashing algorithm to take the inputted name during setup and it is then turned into a multicast address stored in the config file. Thus, once the initial setup is completed, Confluence will use the address this is the reason why user can change the address if needed, without actually changing the name. Consequently the additional nodes using the same multicast address specified in the config file are able to join the cluster.

Each node has a multicast address configured in the `confluence-cfg.xml` file

```
1. name="confluence.cluster.address">xxx.xx.xxx.xxx</property>
```

A warning message is displayed when an user changes the address from the one that Confluence has generated by the hashing of the name. There is no way of eliminating the message any other way other than by returning the address to the one that matches the cluster name. Purpose of the warning message is to remind the user that the address has been changed - as it is not the hashed version any longer - consequently the node can not join the cluster just by using the name. It is also necessary to provide the correct address as well.

Mapping interface to IP address.

To ensure that the interface name is mapped correctly, the following tool can be used. It shows the mapping of the interface name to the IP address.

```
C:\>java -jar list-interfaces.jar
interfaces.size() = 4
networkInterface[0] = name:lo (MS TCP Loopback interface) index: 1 addresses:
/127.0.0.1;

networkInterface[1] = name:eth0 (VMware Virtual Ethernet Adapter for VMnet8) index: 2 addresses:
/192.168.133.1;

networkInterface[2] = name:eth1 (VMware Virtual Ethernet Adapter for VMnet1) index: 3 addresses:
/192.168.68.1;

networkInterface[3] = name:eth2 (Broadcom NetXtreme 57xx Gigabit Controller - Packet Scheduler
Miniport) index: 4 addresses:
/192.168.0.101;
```

Debugging tools

Listed below are some debugging tools that help determine what the status of the multicast traffic is:

Tool	Information provided
`netstat -gn`	Lists multicast groups. Does not work on Mac OS X.
`netstat -rn`	Lists system routing table.
Multicast Test	Coherence tool for testing multicast traffic from one node to another.
`tcpdump -i interface`	Captures network traffic on the given interface. Most useful on an interface that only receives cluster traffic.

Add multicast route

Multicast networking requirements vary across operating systems. Some operating systems require little configuration, while some require the multicast address to be explicitly added to a network interface before Confluence can use it.

If the Multicast Test tool shows that multicast traffic can't be sent or received correctly, adding a route for multicast traffic on the correct interface will often fix the problem. The example below is for a Ubuntu Linux system:

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

To support multiple applications using multicast on different interfaces, you may need to specify a route specific to the Confluence multicast address.

Check firewall

Ensure your firewall allows UDP traffic on the multicast address and port used by Confluence.

Prefer IPv4

> ⛔ There's a known issue with IPv6, especially on Linux.

The fix is to add `-Djava.net.preferIPv4Stack=true` to `JAVA_OPTS`. This tells the JVM to try binding an IPv4 address first, and resort to IPv6 only if that fails.
Note: *A more radical approach is to add NETWORKING_IPV6=no to /etc/sysconfig/network, yet probably should be left for a later consideration on a production machine.*

Change multicast interface

Confluence might have selected the incorrect interface for multicast traffic, which means it cannot connect to other nodes in the cluster. To override the interface used for multicast traffic after initial setup, edit `confluence.cfg.xml` in the Confluence home directory and add a property (or change the existing one) to select your desired network interface. For example to tell Confluence to use `eth1`:

```
1.<property name="confluence.cluster.interface">eth1</property>
```

Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

Create a file in the Confluence home directory called `tangosol-coherence-override.xml`. Add the following to it, setting the TTL value appropriately (1 is the default):

```
1.<?xml version='1.0'?>
2.<coherence>
3.<cluster-config>
4.    <multicast-listener>
5.        <time-to-live system-property='tangosol.coherence.ttl'>1</time-to-live>
6.    </multicast-listener>
7.</cluster-config>
8.</coherence>
```

Alternatively, simply start Confluence with the system property: `-Dtangosol.coherence.ttl=1`. Again, 1 is the default value, and you should change it to something appropriate to your network topology.

Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain operating systems.

For best results, use the simplest network topology possible for the cluster traffic between the nodes. For two nodes, that means a single network cable. For larger numbers, try using a single high-quality switch.

Advanced Tangosol configuration

If the solution to your problem involves changes to the Tangosol configuration, these changes should **not** be made to the Confluence configuration in `confluence/WEB-INF/classes/`. Instead, to ensure your configuration survives upgrades, make your changes via:

* Tangosol system properties
* creating a `tangosol-coherence-override.xml` file in the Confluence home directory.

Examples of making these changes are shown in the increasing the TTL section.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

How do I supress cluster warning message in confluence?

Recommended network topology

Apache and Tomcat load balancing

Cluster Panic triggers

Confluence Cluster Installation

Confluence Clustering Overview

Changing Datasources Manually in a Cluster

Technical Overview of Clustering in Confluence

Viewing and Editing License Details

Cluster safety mechanism

Upgrading a Confluence Cluster

Cluster Troubleshooting

Cluster Administration page

Open JIRA Features and Bug Reports

JIRA Issues (44 issues)								
Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Update
		Attachment migration does not		Nicholas					

	CONF-8959	happen when upgrading to a clustered license	Unassigned	Ilacqua [Atlassian]	⬆	Open	Unresolved	Jul 19, 2007	Jul 02, 2009
	CONF-9297	Confluence should be able to automatically recover from cluster panics	Unassigned	Gary Weaver	⬆	Open	Unresolved	Aug 27, 2007	Mar 25, 2009
	CONF-12689	Support Confluence cluster upgrades without an outage	Unassigned	Igor Minar	⬆	Open	Unresolved	Aug 06, 2008	Sep 10, 2008
	CONF-14120	Hibernates UpdateTimestampsCache doesn't handle concurrent writes	Unassigned	Chris Kiehl [Atlassian]	⬆	Open	Unresolved	Jan 05, 2009	May 05 2009
	CONF-12287	Coherence cache fails while retrieving profile picture metadata (dashboard or view page shows UnexpectedRollbackException)	Unassigned	Matt Ryall [Atlassian]	⬆	Open	Unresolved	Jul 01, 2008	Dec 08, 2009
	CONF-8716	Determine index mismatch in cluster and warn on cluster info page	Unassigned	Matt Ryall [Atlassian]	⬆	Open	Unresolved	Jun 17, 2007	Oct 29, 2007
	CONF-9040	Authenticator (subclass of DefaultAuthenticator) can be called twice at almost exactly same time by 2 or more clustered servers	Unassigned	Gary Weaver	⬆	Open	Unresolved	Jul 30, 2007	Nov 04, 2007
	CONF-9594	ConditionalPropertySet's cannot be cached breaking cluster installations that delegate user management to JIRA	Unassigned	Dave Loeng [Atlassian]	⬆	Open	Unresolved	Sep 28, 2007	Jul 02, 2009
	CONF-9813	Disable attachments migration to Filesystem in Cluster	Unassigned	Gurleen Anand [Atlassian]	⬆	Open	Unresolved	Oct 24, 2007	Oct 25, 2007
	CONF-10054	OSUser provider does not generate IDs correctly in clustered environment	Unassigned	Christopher Owen [Atlassian]	⬆	Open	Unresolved	Nov 26, 2007	Dec 01, 2008
	CONF-10323	Coherence Lock being held when it appears no thread should have the lock. Causes ConcurrentModificationException	Unassigned	Paul Curren [Atlassian]	⬆	Open	Unresolved	Dec 26, 2007	Sep 17, 2009
	CONF-10325	Viewing the members of a group in a clustered environment works only on one node and not the other.	Unassigned	Partha Kamal [Atlassian]	⬆	Open	Unresolved	Dec 27, 2007	Jul 02, 2009
	CONF-10868	Node that can not join cluster due to license restriction causes cluster panic	Unassigned	Ivan Benko [Atlassian]	⬆	Open	Unresolved	Feb 29, 2008	Sep 03, 2008
	CONF-10980	Cluster debugging/troubleshooting tools	Unassigned	Ivan Benko [Atlassian]	⬆	Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-12345	Park issue :)	Unassigned	Dave Loeng [Atlassian]	⬆	Open	Unresolved	Jul 07, 2008	Jan 21, 2009
	CONF-12486	ClassNotFoundException logged on cluster node startup	Unassigned	Anatoli Kazatchkov [Atlassian]	⬆	Open	Unresolved	Jul 17, 2008	Aug 25, 2009
	CONF-12614	Intermittent ConcurrentModificationException in cluster	Unassigned	Anatoli Kazatchkov [Atlassian]	⬆	Open	Unresolved	Jul 29, 2008	Mar 31, 2009
	CONF-13421	Layout customisations are not propagated to other cluster nodes	Unassigned	Matt Ryall [Atlassian]	⬆	Open	Unresolved	Oct 16, 2008	Dec 09, 2008
	CONF-14045	Cluster mysql tests corrupts the database	Unassigned	Brian Nguyen [Atlassian]	⬆	Open	Unresolved	Dec 21, 2008	Jan 08, 2009
	CONF-14076	SchedulerException when running cluster builds	Unassigned	Brian Nguyen [Atlassian]	⬆	Open	Unresolved	Dec 29, 2008	Jan 01, 2009
	CONF-14657	Retrieving the global settings in a clustered environment causes a lot of contention	Unassigned	Chris Kiehl [Atlassian]	⬆	Open	Unresolved	Feb 21, 2009	Nov 08, 2009
	CONF-14948	Support failover NICs for cluster configuration...	Unassigned	Tony Atkins [Atlassian]	⬆	Open	Unresolved	Mar 19, 2009	Mar 19, 2009

	CONF-16419	Installing a font for PDF export in a cluster will not carry to cluster nodes that are down or unavailable.	Unassigned	Charles Miller [Atlassian]	⬆	Open	Unresolved	Jul 20, 2009	Aug 05, 2009
	CONF-9335	In cluster, allow attachments to be stored on file system in network-shared directory	Unassigned	Jeremy Largman [Atlassian]	⬆	Open	Unresolved	Aug 29, 2007	Nov 29, 2009
	CONF-7368	Confluence Cluster setup dies horribly when DNS fails	Unassigned	Don Willis [Atlassian]	⬇	Open	Unresolved	Nov 23, 2006	Jan 19, 2009
	CONF-9020	Cluster nodes do not get notified of Layout changes	Unassigned	Roberto Dominguez	⬇	Open	Unresolved	Jul 27, 2007	Oct 08, 2007
	CONF-9059	Cluster setup network interface selection shows loopback interface	Unassigned	Matt Ryall [Atlassian]	⬇	Open	Unresolved	Aug 01, 2007	Aug 01, 2007
	CONF-9281	Plugin's I18n properties not loaded in other cluster nodes unless restarted	Unassigned	Roberto Dominguez	⬇	Open	Unresolved	Aug 26, 2007	Sep 15, 2008
	CONF-9749	Coherence does not allow the disabling of all JDK shutdown hooks	Unassigned	Christopher Owen [Atlassian]	⬇	Open	Unresolved	Oct 17, 2007	Jan 29, 2008
	CONF-10635	Database logging of clustersafety access	Unassigned	James Fleming [Atlassian]	⬇	Open	Unresolved	Feb 03, 2008	Feb 03, 2008
	CONF-10953	Support unicast addressing in cluster when well-known-addresses WKA are defined	Unassigned	Ivan Benko [Atlassian]	⬇	Open	Unresolved	Mar 06, 2008	Aug 25, 2009
	CONF-10977	Generate new Multicast address from a "new" cluster name	Unassigned	Ivan Benko [Atlassian]	⬇	Open	Unresolved	Mar 06, 2008	Sep 11, 2008
	CONF-10979	List confluence cluster interface	Unassigned	Ivan Benko [Atlassian]	⬇	Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-10981	Check how many nodes/processes running in a cluster and their identity	Unassigned	Ivan Benko [Atlassian]	⬇	Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-12421	Don't use distrbuted cache for storing Captcha's in a cluster	Unassigned	Paul Curren [Atlassian]	⬇	Open	Unresolved	Jul 13, 2008	Jul 13, 2008
	CONF-13698	Changing custom html on one node of a cluster is not imideatly reflected on the other node.	Unassigned	Anatoli Kazatchkov [Atlassian]	⬇	Open	Unresolved	Nov 12, 2008	Nov 13, 2008
	CONF-13870	After a site Import into a cluster, admin console displays attachment storage as filesystem	Unassigned	Agnes Ro [Atlassian]	⬇	Open	Unresolved	Nov 27, 2008	Dec 03, 2008
	CONF-14338	Specify an arbitrary multicast port for a cluster	Unassigned	James Fleming [Atlassian]	⬇	Open	Unresolved	Jan 30, 2009	Jan 30, 2009
	CONF-17040	Cannot build milestones outside Atlassian due to coherence	Unassigned	Jonathan Gilbert [Atlassian]	⬇	Open	Unresolved	Sep 25, 2009	Dec 08, 2009
	CONF-9712	Plugins which don't work in a cluster shouldn't look like an error	Unassigned	Gary Weaver	⬇	Open	Unresolved	Oct 15, 2007	Oct 16, 2007
	CONF-9846	Tangosol configuration: the (optional) cluster-name element is in the wrong place	Unassigned	Don Willis [Atlassian]	⬇	Open	Unresolved	Oct 30, 2007	Dec 20, 2007
	CONF-14088	Locking on cache keys needs to check if the lock was actually aquired	Unassigned	Chris Kiehl [Atlassian]	⬇	Open	Unresolved	Dec 30, 2008	Nov 22, 2009
	CONF-10330	Certain Multicast IP addresses sending CHANGE_TO_EXCLUDE_MODE causing IGMP traffic to be blocked	Unassigned	Partha Kamal [Atlassian]	⬆	Needs Verification	Unresolved	Dec 28, 2007	Nov 26, 2008
	CONF-8300	Cannot override TTL configuration through tangosol coherence properties	Unassigned	Matthew Jensen [Atlassian]	⬆	Needs Verification	Unresolved	Apr 20, 2007	Nov 10, 2009

Contact Atlassian support

We have dedicated staff on hand to support your installation of Confluence. Please follow the instructions for raising a support request and mention that you're having trouble setting up your Confluence cluster.

Related

Cluster Safety Mechanism

Cluster Panic triggers

Cluster Panic

Usually presents itself with the following error message:

```
FATAL [DefaultQuartzScheduler_Worker-4] [confluence.cluster.safety.ClusterPanicListener]
handleEvent Fatal error in Confluence cluster:
Database is being updated by an instance which is not part of the current cluster. You should
check network connections between cluster nodes, especially multicast traffic.
```

In almost all cases, cluster panic events are caused by two or more instances of Confluence (in separate clusters) updating the same database. Such events are typically caused by one of the following issues:

JVM paused (e.g. while swapping memory) can break communication between two nodes

Always watch the swapping activity of your server and avoid swapping due to lack of RAM. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection (GC) considerably and affect Confluence's performance.

- In clustered installations, swapping can lead to cluster panic. This is because swapping causes the JVM to pause during garbage collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Two instances of Confluence have been started in your application server

This is one of the most commonly encountered issues. The strangest case of this that we have seen so far involved a cloned image of a PC running Confluence that was later used in a remote office in a different city. The people using Confluence on the cloned instance were not aware that the original Confluence instance was also running and that both these Confluence instances were using the same production database server.

- Solution: Check your application server's configuration to make sure that multiple copies of the application server are not running concurrently. Database transaction logs can help identify the location of other application servers, if client IP addresses are recorded along with each transaction.

Two copies of your application server are running.

Sometimes starting an application server twice will result in two processes running, even though only one can be accessed over the network.

- Solution: Check a list of running processes (for example, with the 'ps' command in Posix-based operating systems like Linux, Unix and Mac OS X) and make sure your application server is only running once.

Networking failure between nodes in the cluster

- Solution: Check that multi-cast traffic is being transmitted successfully, and that the network between your nodes is low-latency (<100 ms).

Database server stops responding

If Coherence fails to retrieve the SafetyNumber from the database, the comparison will fail. If it fails to update it, the next comparison will fail, 30 seconds later.
Many things can cause this, including a scheduled shutdown for backups, network failure, a filled-up transaction-log partition and a changed password on the account used by Confluence to connect to the database.

- Solution: resolve the problem with the database (or network), then restart Confluence

> 🚫 In all cases, when starting Confluence after a cluster panic, you must ensure all cluster nodes have been shut down completely. If necessary, use commands like `ps` and `kill` to get a list of Java processes and terminate them manually.

> ⚠️ Please visit this document for troubleshooting advice if you encounter any of the above situations.

Multicast Test

This page describes the **Multicast Test**, a Coherence tool for testing multicast traffic from one node to another. You may find this useful when troubleshooting a clustered installation of Confluence.

In order to run the Multicast test, you need to first download the attached Coherence zip file.

The Multicast Test comes as a script called `multicast-test`, which you will find located in the `bin` folder in the above zip file.

Instructions on how to run this script file can be found in the Coherence documentation. You may like to go straight to the subheading called '**Example**' in the guide, where there is an example on how to use the `multicast-test` script.

RELATED TOPICS

Cluster Troubleshooting
Confluence Clustering Overview

Clustering for Scalability vs Clustering for High Availability (HA)

People occasionally enquire about setting up High-Availability (HA) Confluence clusters. Confluence's clustering is designed to solve a different problem, that of scaling under high load. This page explains the difference.

On this page:

- What is High Availability (HA)?
- What does Confluence's clustering do, then?
- So what kind of resilience can I build into a Confluence installation?
- What's the difference between load balancing and failover?
- What do you mean by 'session affinity'?

What is High Availability (HA)?

HA means that your application *will be* available, without interruption. It's a very difficult thing to achieve, and is typically what people are talking about when they refer to five-nines availability.

In the context of application clustering, it means that any given node (or combination of nodes) can be shut down, blown up, or simply disconnected from the network unexpectedly, and the rest of the cluster will continue operating cleanly as long as at least one node remains. It requires that nodes can be upgraded individually while the rest of the cluster operates, and that no disruption will result when a noes rejoins the cluster. It typically also requires that nodes be installed in geographically separate locations.

What does Confluence's clustering do, then?

Confluence's clustering system allows a single installation to serve a much greater number of concurrent requests than a single server. This is what we refer to as 'scaling under load'.

It does provide a certain amount of resilience, as the death of one node won't bring the other(s) down. However, it requires very low network latency, which rules out geographic separation of the servers, and upgrading can only be performed while the entire cluster is shut down. This doesn't mean that Confluence's clustering is buggy or broken. It simply reflects the difference between the two design aims.

So what kind of resilience can I build into a Confluence installation?

It's still entirely possible to build a resilient Confluence installation, using a 'cold-failover' approach in which two (or more) servers share a database and (normally) a network-mounted file system, where no more than one server is actually running at any given time.

Several different approaches are feasible, but the common elements are:

- a well-configured load balancer (session affinity is irrelevant in this case)
- a reliable monitoring system which can detect and shut down a misbehaving Confluence instance before starting the spare server
- startup scripts with added smarts to check for the presence of another running node before deciding whether to start up a server
- servers with the same view of both the database and the home directory.

> ⚠ It's vital to ensure that only one server is running at any one time, in this kind of setup. If a server starts while another is already running against the same database, the result will be a cluster panic that shuts down both servers.

A single database becomes the single point of failure in such a system. This can be alleviated by database clustering, or by replication from the 'active' database server to the standby server(s) if you wish to separate the failover systems while keeping database latency to a minimum.

In the same vein, the home directory can be hosted on a shared network system — SAN or NAS, preferably with its own replication/rapid recovery system — though there's a known issue to consider. Alternatively, to avoid the use of networked file systems, a utility such as `rsync` can be used to periodically bring the spare servers' home directories up to date, so long as you keep the period sufficiently short — probably between one and five minutes, depending on the rate of activity. This can be avoided altogether by keeping attachments in the database; it increases the demands on the bandwidth between the application and database servers, but guarantees that the system is in a consistent state at switchover. If the data is at all sensitive or confidential, it's advisable to run `rsync` over `ssh`, to minimise the opportunity for the data to be captured on its way across the network.

What's the difference between load balancing and failover?

Load balancing means that all servers are active, and new requests are distributed among them. Several strategies are available, but the most common are:

- round-robin — the first request goes to the first server, the second request goes to the second server, and so on. When you run out of servers, the next request goes to the first server, and around it goes again.
- percentage-based — if (for example) you have two servers, and one can handle twice the load of the other, you can tell the load balancer to send two requests to the stronger server for every request that goes to the weaker one.
- availability — the load balancer sends a test query to each of the servers every second or so, and directs each new request to the server that's currently responding the fastest.

Failover means that only one server is active at any given time, and normally involves two servers (any number of servers may be involved, depending on the system). If the active one stops responding, requests are directed to the other server — the system 'fails over' to the second one.

'Cold failover' means that the second server is only started up after the first one has been shut down. This is the case for non-clustered Confluence.

'Hot failover' or 'hot standby' means that all servers are running at all times, and that the load is directed entirely toward one server at any one time.

A load balancer can be used in both scenarios, especially if it's smart enough to keep track of which servers are currently running.

Failover can also be managed via DNS, in a sufficiently well-controlled environment.

What do you mean by 'session affinity'?

Sessions consist of several transmissions in each direction between the client (browser) and the server. Session affinity means that the load balancer keeps track of which server received the initial transmission from a given browser, and that it will then send any subsequent requests from that browser to the same server.

This is necessary with Confluence clustering, in particular, because sessions are not shared across cluster nodes. If you log into one node and then send a request to another, the other node will send you the login screen because it doesn't recognise your session cookie.

RELATED TOPICS

Confluence Clustering Overview

Recommended network topology

Atlassian recommends a network topology similar to the one shown below, to get the best results from a Confluence Clustered deployment.

The number of Confluence nodes in the deployment is adjustable — select the number which suits your own requirements.

The most important aspect is that cluster, database and HTTP (client) traffic are all carried on separate subnets. It is possible, on a sufficiently fast network, to carry cluster and database traffic on the same subnet but we do strongly recommend that HTTP traffic be always confined to a separate subnet on production deployments.

Confluence Clustered does not support clustered communication over WAN, VLAN or VPN. All Confluence Clustered nodes must be on the same local subnet, ideally networked via an ethernet hub or simple switch. The cluster communication network must also support multicast IP networking.

> ✅ **Use this example as a basis for your own network diagram**
> When you are considering a Confluence Clustered deployment, you should prepare a network diagram like the one on this page. This will facilitate discussion with Atlassian Support and help with your own planning. Please refer to the cluster checklist for more guidance on planning your clustered deployment.
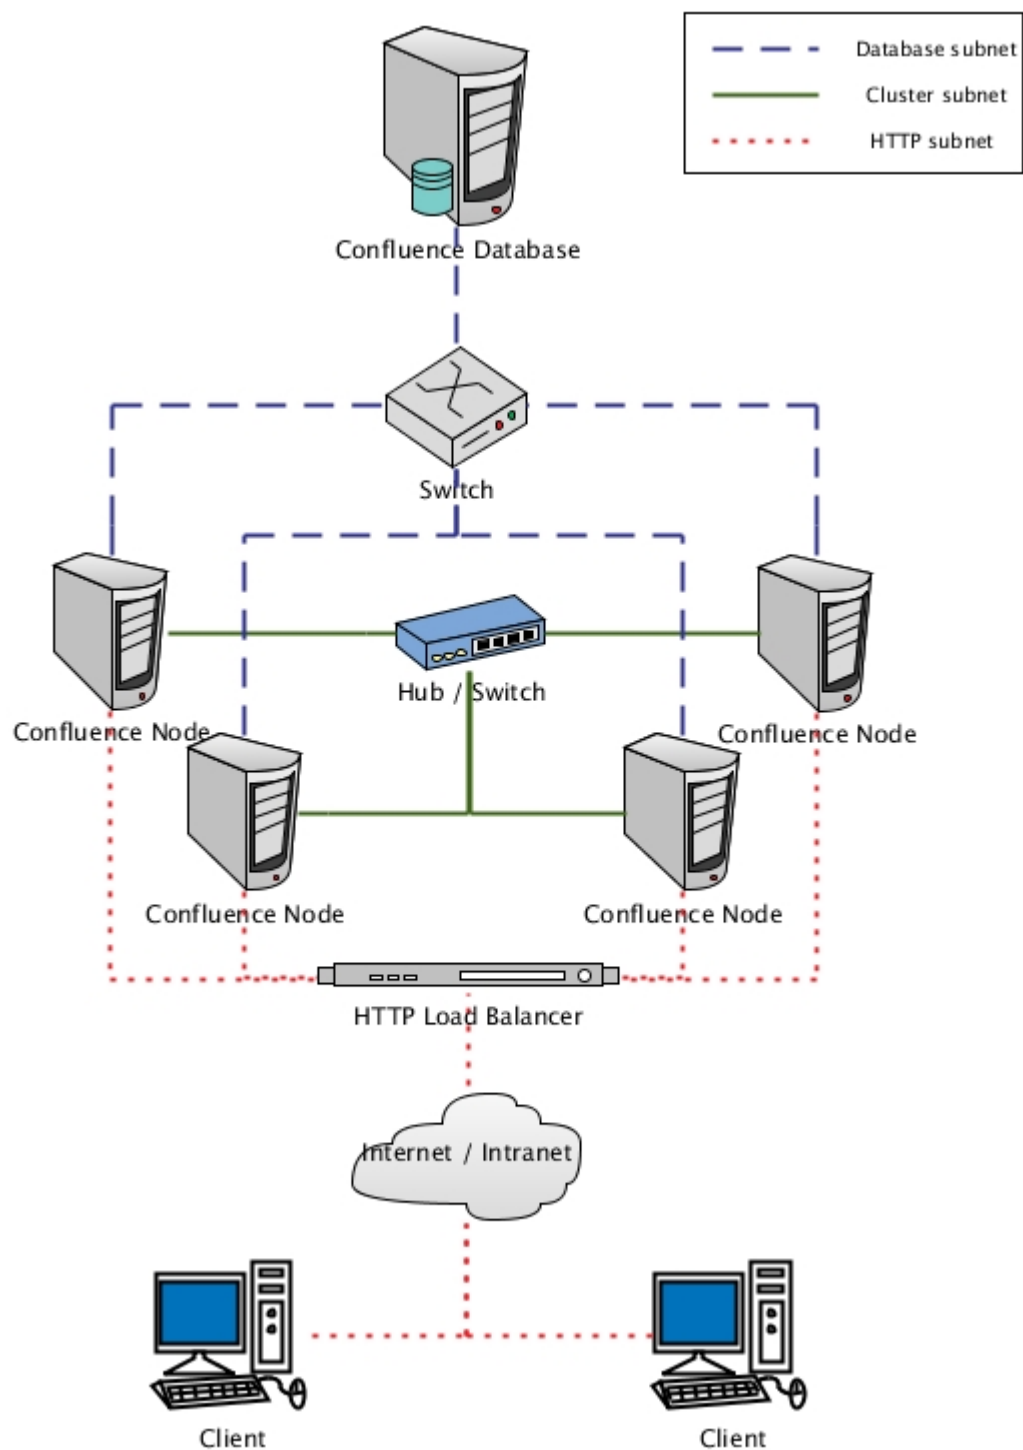
Cluster Administration page

> ℹ **Overview of clustering documentation**
> Refer to the overview of Confluence clustering.

Overview

Any instance of Confluence which uses a clustered license has a Cluster Configuration page which includes information about the active cluster.

To open the Cluster Administration page,

> 1. Go to the Confluence '**Administration Console**'. To do this:
>
>   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
> 2. Click '**Cluster Configuration**' in the left-hand menu, in the section called 'Clustering'.

Availability

To access this functionality, you must:

- Be a System Administrator (i.e. have global System Administrator permissions), and
- be using Confluence 2.3 or later, and
- be using a clustered Confluence license.

Screenshot: Cluster Administration Page

This page shows your cluster configuration, and allows you to start a new Confluence cluster using data from this instance.

Cluster Status indicates whether your cluster is currently running.

Licensed nodes is the maximum number of instances of Confluence your license allows in a cluster.

Active nodes lists the instances of Confluence currently participating in the cluster.

Starting a new cluster will perform the following changes:

- enable a clustered cache
- migrate attachments from file system to the database
- publish database connection information so other nodes can join the cluster.

ℹ All access to Confluence will be locked while this takes place, and you will be forced to restart Confluence afterwards.

Cluster name is a short name for identifying your cluster. Other Confluence instances can join the cluster using this name.

🛈 To join an existing cluster, start a clean copy of Confluence on this node and select 'Join Cluster' during the setup wizard.

Related documents

Overview of Confluence Clusters
Confluence Cluster Installation
Cluster Troubleshooting

Cluster Checklist

It is possible to run Confluence in a **clustered** environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.

🛈 Refer to the clustering overview for more information and a list of related pages about clustering Confluence.

⚠️ **Consider your options carefully before deciding on a clustered installation**
While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information below and then consult Atlassian Sales before making your final decision.

Summary of the information on this page:

- Purpose of this Document
- Assumed Knowledge
- General Considerations
  - Confluence Clustered is designed to scale the number of simultaneously connected users at a much better performance than what a single node can achieve
  - Confluence Clustered will not improve performance in systems with few users.
  - Confluence Clustered is not a high availability solution.
  - Confluence Clustered is not for disaster recovery nor for transparent failover.
- Server Setup
  - The number of supported cluster nodes is limited to four.
  - All cluster nodes must have the same version of OS, application server, etc.
  - Use good and up-to-date hardware.
  - Confluence Clustered is not supported when run in VMware or other virtualisations.
  - Confluence should be the only application on the cluster servers.
  - Do not upgrade and switch to Confluence Clustered at the same time
- Database Setup
  - Run the database on its own physical server.
  - Attachments must be stored in a database and not the local file system
  - Make sure that you use a supported version of a database server to store Confluence's data.
  - Your database must be provisioned to store a large volume of binary data.
  - You need an experienced DBA available to troubleshoot database performance issues.
- Network Setup
  - We recommend hardware load balancers or putting a software loadbalancer onto its on server.
  - Use separate network adapters for communication between servers.
  - The switch connecting the Confluence cluster nodes must not be a 'smart switch'.
  - Cisco switches need additional configuration.
  - It is recommended that the database is on a different physical network from the Confluence server nodes.
  - Minimize the latency between the Confluence cluster nodes and the database.
  - Prepare a network diagram.
  - You need network support staff available to troubleshoot cluster communication issues.
- Staging Environment
  - You need a staging environment that is exactly the same as your production system.

Purpose of this Document

The purpose of this cluster checklist is to help you:

- Decide whether Confluence Clustered is the right solution for you.
- Create a plan for your clustered deployment.

As a service to our customers, we offer to review your deployment plan and make recommendations to help you avoid common pitfalls. To make use of this service, please consider all the information below carefully while planning your clustered deployment. Then contact Atlassian Pre-Sales for recommendations.

If you need to raise a support request with Atlassian during or after cluster deployment, we will need to ask you questions about your configuration. It will save crucial time if you can provide us with your deployment plan.

For more information about clustering Confluence, refer to the clustering overview.

Assumed Knowledge

In writing this document, we have assumed that our readers have an in-depth knowledge of the following technical areas:

- Database
- Networking
- Application servers
- Load balancers

Before starting a clustered deployment please read the information on this page carefully, as well as the linked documentation, to assess if you have the assumed knowledge.

General Considerations

> **What will Confluence Clustered do for you?**
> The points in this section of the page will help you evaluate your reasons for considering a clustered deployment, and then decide whether Confluence Clustered is the right solution for your environment.

Confluence Clustered is designed to scale the number of simultaneously connected users at a much better performance than what a single node can achieve

Confluence Clustered will not improve performance in systems with few users.

Clustering Confluence means that user requests can be served by independent machines. The performance gains are substantial, and have improved a lot further since Confluence 3.0. Clustering is especially great in dealing with spikes to the load, e.g. during certain hours of business. Just note that if rendering a complicated page (e.g. containing many macros or rendering many graphs) takes five seconds on an otherwise idle server, it will not be faster in a clustered environment. Also, the first step when you encounter performance issues is to tune your existing system, make sure you are using the right hardware and have looked at your database.

Confluence Clustered is not a high availability solution.

Confluence Clustered is not designed specifically to provide a high availability solution.

General availability is higher in a Confluence cluster than on a single installation, you can for example take one node down for minor maintenance tasks e.g. when adding a new CPU or adding RAM. But you still have to bring down all nodes at the same time for software upgrades. Also there are certain conditions, like loss of network connectivity between nodes ('split brain'), that will result in the cluster shutting itself down. Confluence Clustered offers higher reliability, but not high availability.

Confluence Clustered is not for disaster recovery nor for transparent failover.

If one node crashes, there is no transparent failover for the connected client. Also, our network requirements (see below) make Confluence unsuitable for deployment to different cities or even to different buildings.

Server Setup

The number of supported cluster nodes is limited to four.

⚠ **Not supported.** In theory, you can connect more than four nodes — but that is not covered by Atlassian Support.

All cluster nodes must have the same version of OS, application server, etc.

Confluence requires a homogeneous environment. All Confluence cluster nodes must have the same version of the following:

- Operating system
- CPU
- Installed memory
- Java
- Application server

ℹ Note that 'same version' means '**same to the last digit**'. For example, Java v1.4.2_16 is not the same as v1.4.2_15

✅ We strongly recommend user to have the same memory configuration (both the JVM and the physical memory) because a cluster uses a replicated cache. A replicated cache requires the same amount of memory on each node in the operating cluster. The memory allocations must be equal.

Use good and up-to-date hardware.

While the details are up to you, we strongly suggest that your servers have at least 4GB of physical RAM. A high number of concurrent users

means that a lot of RAM will be consumed. You usually don't need to assign more than 4GB per JVM process, and most of the time even just 1GB or 2GB will be fine, you should just be prepared to fine tune the settings.

Confluence Clustered is not supported when run in VMware or other virtualisations.

⚠ **Not supported.** We strongly discourage you to deploy a production environment of Confluence to virtual servers, and we will not be able to support you when problems arise.

When running a Confluence cluster your goal is high capacity and performance, so you should not risk lower performance by virtualising it and sharing a computer with other processes.

Many customers who are running Confluence on VMware, or similar virtualisation solutions, experience major performance problems that are extremely hard to pinpoint. Since the problems are not related to Confluence itself, we will not be able to help you.

Confluence should be the only application on the cluster servers.

No additional applications (other than core operating system services) should be running on the same servers as Confluence.

Since your goal should be increased capacity and performance, you should not risk this by running any other process on the machine with a Confluence Clustered node. While it may be fine to run JIRA, Confluence and Bamboo on a dedicated Atlassian software server for small installations, it is strongly discouraged for clustering Confluence.

Do not upgrade and switch to Confluence Clustered at the same time

If you plan to migrate to a clustered solution, make sure you are migrating within the same version of Confluence. If you plan to upgrade to a higher version of Confluence, do this **before** the migration to the clustered version.
For example, if you are currently running Confluence 2.9.2 standalone, and want to roll out the clustered version of Confluence 3.0, you must first upgrade to Confluence 3.0 standalone and check that everything works fine (e.g. by running and monitoring your production system for a week). Then you are in a good position to migrate to the clustered version.

Database Setup

Run the database on its own physical server.

You are optimising for performance, so you don't want the database to slow down your application servers, or vice versa. In high load scenarios, the database may need to have better hardware than the application servers to be able to handle all requests. You should find out by performing loadtesting.

Attachments must be stored in a database and not the local file system

Storing attachments in the database is the only supported attachment storage configuration for clustering Confluence.

Make sure that you use a supported version of a database server to store Confluence's data.

Please check that your intended database is officially supported by Atlassian Confluence. The load on an average cluster solution is higher than on a single box installation, and it is therefore even more crucial to use the right database vendor and version.

Your database must be provisioned to store a large volume of binary data.

Note that Confluence clustered stores file attachments in the database, and you need an experienced DBA who can monitor and manage the data growth.

You need an experienced DBA available to troubleshoot database performance issues.

Not having an experienced full-time DBA at hand at short notice when entering the realm of high load is dangerous. While small installations of Confluence basically work 'out of the box', anything that involves high load and a lot of database space requires continual monitoring, optimising and fine tuning of the Confluence database. When we ramp up the load on our loadtesting environment, we see that database usage goes up as well. Having powerful hardware in place helps, but if there are queries that become inefficient with you particular load pattern, you need an expert to tune it. As an example, we have seen PostgresSQL switch its internal caching mechanism when a particular table reached a certain size, which resulted in a drop of performance by about 200ms per request. This happened from one second to the other. Being able to troubleshoot and then fix issues like these is important in any enterprise system, but it is even more in a high load scenario.

Network Setup

We recommend hardware load balancers or putting a software loadbalancer onto its on server.

If you use a software load balancer (which is fine except for really extreme installations), it must be deployed on a machine of its own. Running a software load balancer on a cluster node is not supported. If a node unexpectedly got overwhelmed by a spike in load, a load balancer on that node would turn unresponsive. As a result, your whole cluster would be inaccessible even though the other nodes would be available. So using a different server is common practice and common sense.

Use separate network adapters for communication between servers.

The Confluence cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication.

This is the best way of getting the cluster to run fast and reliably. Performance problems are likely to occur if you connect cluster nodes via a network that has lots of other data streaming through it.

The switch connecting the Confluence cluster nodes must not be a 'smart switch'.

⚠ **Not supported.** Smart switches are not covered by Atlassian Support for Confluence Clustered.

Do not use smart switches between cluster nodes. Many problems have been reported and attributed to smart switches. They have a tendency to interrupt broadcast or multicast traffic, thus reliably killing a cluster after a certain amount of time has passed. This makes troubleshooting especially complex and tedious.

Cisco switches need additional configuration.

If the switch connecting the Confluence cluster nodes is a Cisco switch then it might need additional configuration to support Confluence clustering.

Please make sure you find out all the details about your switches before you start the deployment.

It is recommended that the database is on a different physical network from the Confluence server nodes.

Since you want to increase your capacity and performance for high loads, it is recommended to have your database on a different network. Please refer to the recommended topology diagram for more information.

Minimize the latency between the Confluence cluster nodes and the database.

Even though having the nodes and the database on the same physical network usually suffices, you should take the time to explicitly measure network latency, and make sure it is as close to zero as possible.

Prepare a network diagram.

To facilitate discussion and to ease planning, you should prepare a network diagram like this example of recommended network topology.

If you request support with Confluence Clustered, we may ask for your network diagram. We recommend that you create one similar to our example before you proceed with the installation.

You need network support staff available to troubleshoot cluster communication issues.

Setting up a cluster is not trivial. Even small problems in network design will be expanded in a clustered installation. (This is true of any kind of software.)

It is absolutely vital that you have dedicated network staff available to track down problems when they arise. A cluster will usually be used by thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone made an undocumented change to the network and you don't have an expert around who can figure it out.

Staging Environment

You need a staging environment that is exactly the same as your production system.

You must be able to test drive any change to the cluster (installing upgrades, installing plugins) and to perform other tests (checking connectivity, debugging problems) on a staging cluster.

The staging environment must be:

- On the same OS, database, and Java version as your production environment.
- Clustered.

If you require support, we may for example ask you to turn off certain third-party plugins. If you can't do this in your production environment and you don't have a staging environment for troubleshooting, we may not be able to help you.

> 🛈 **Getting a license for your staging environment**
>
> > 🛈 **Only a technical contact for your commercial/academic license is able to create a Developer license**
>
> Atlassian supplies 'developer' licenses which can be used by existing commercial license holders who wish to deploy non-production installations of our software to use in QA/staging environments. Developer licenses are free of charge to commercial license holders and, like our commercial offerings, they include 12 months of updates starting from the date of purchase of the commercial license.
>
> If you hold a commercial license, you can obtain a free developer license by performing the following:
>
> 1. Log in to your Atlassian account.
> 2. Under the "Licenses" heading, all of your licenses will be displayed. Click the plus sign next to a license to view its details.
> 3. Click the **'View Developer License'** link in the bottom right corner of the license detail panel, below your commercial license key.

RELATED TOPICS

How do I supress cluster warning message in confluence?

Recommended network topology

Apache and Tomcat load balancing

Cluster Panic triggers

Confluence Cluster Installation

Confluence Clustering Overview

Changing Datasources Manually in a Cluster

Technical Overview of Clustering in Confluence

Viewing and Editing License Details

Cluster safety mechanism

Upgrading a Confluence Cluster

Cluster Troubleshooting

Cluster Administration page

Confluence Security

As a public-facing web application, Confluence's application-level security is obviously important. This document answers a number of questions that commonly arise when customers ask us about the security of our product.

> 🛈 This document is for system administrators looking to evaluate the security of the Confluence web application. It does not address Confluence's internal security model (user/group management and content permissions), except as it relates to the overall application security. For information about user management, groups and permissions, please refer to the internal security overview.

On this page:

- Application Security Overview
  - Password Storage
  - Buffer Overflows
  - SQL Injection
  - Script Injection
  - Cross-Site Scripting
  - Transport Layer Security

Application Security Overview

Password Storage

When Confluence's internal user management is used, passwords are hashed through SHA1 before being stored in the database. There is no mechanism within Confluence to retrieve a user's password – when password recovery is performed, a new random password is generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities – possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customisations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavour to fix them as quickly as possible.

Transport Layer Security

Confluence does not directly support SSL/TLS. Administrators who are concerned about transport-layer security should set up SSL/TLS at the level of the Java web application server, or the HTTP proxy in front of the Confluence application.

For more information on configuring Confluence for SSL, see: Adding SSL for Secure Logins and Page Security

Session Management

Confluence delegates session management to the Java application server in which it is deployed. We are not aware of any viable session-hijacking attacks against the Tomcat application server shipped with Confluence Standalone. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

Administrator Trust Model

Confluence is written under the assumption that anyone given System Administrator privileges is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged

network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a `chroot` jail.

Stack Traces

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in How to Report a Security Issue.

Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in Security Advisory Publishing Policy.

Severity Levels

Atlassian's approach to ranking security issues is detailed in Severity Levels for Security Issues.

Our Patch Policy

Atlassian's approach to releasing patches for security issues is detailed in Security Patch Policy.

Published Security Advisories

- Confluence Community Security Advisory 2006-01-19
- Confluence Security Advisory 2005-02-09
- Confluence Security Advisory 2005-12-05
- Confluence Security Advisory 2006-01-20
- Confluence Security Advisory 2006-01-23
- Confluence Security Advisory 2006-06-14
- Confluence Security Advisory 2007-07-26
- Confluence Security Advisory 2007-08-08
- Confluence Security Advisory 2007-11-19
- Confluence Security Advisory 2007-11-27
- Confluence Security Advisory 2007-12-14
- Confluence Security Advisory 2008-01-24
- Confluence Security Advisory 2008-03-06
- Confluence Security Advisory 2008-03-19
- Confluence Security Advisory 2008-05-21
- Confluence Security Advisory 2008-07-03
- Confluence Security Advisory 2008-09-08
- Confluence Security Advisory 2008-10-14
- Confluence Security Advisory 2008-12-03
- Confluence Security Advisory 2009-01-07
- Confluence Security Advisory 2009-02-18
- Confluence Security Advisory 2009-04-15
- Confluence Security Advisory 2009-06-01
- Confluence Security Advisory 2009-06-16
- Confluence Security Advisory 2009-08-20
- Confluence Security Advisory 2009-10-06
- Confluence Security Advisory 2009-12-08

Related Server Security Pages

Adding SSL for Secure Logins and Page Security

▷ Click to see pages related to user and group permissions.

Assigning Space Permissions

Confluence Cookies

Confluence Security

View File Macro

Revoking Space Permissions

Edit in Word Link Macro

Hiding the People Directory

Confluence Security Advisory 2006-01-23

Security Overview

HTML Macro

How to Hide the Referrer

Space Permissions Overview

Confluence Community Security Advisory 2006-01-19

> ⛔ This security advisory is not endorsed by Atlassian - this is a public service advisory from a member of the confluence community. **Please** remember to backup any modified files, and use these instructions at your own risk. While this information is based on Confluence v2.1.2, it may have uses with older affected versions of Confluence.
>
> **The official security advisory is located at Confluence Security Advisory 2006-01-20**

Problem

There is a possibility of XSS exploitation of the Full Name user profile field when displayed.

Solution

The problem was unescaped outputting of the fullname - wrapping the output in $generalUtil.htmlEncode() resolve it. The vast majority of the problem can be resolved by changing `/confluence/template/includes/macros.vm` in the distribution on the following lines:

- 180
- 186
- 200
- 340
- 893

I have attached the modified macros.vm file here which you can copy into your distribution.

Scope

There are other places which are still affected which Atlassian have been made aware of, a complete resolution should be provided by Atlassian in their own offical advisory.

I hope this helps some of you!

Confluence Security Advisory 2005-02-09

A flaw has been found in Confluence by which attackers can bypass Confluence security and change content on the site. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 1.3.3

Vulnerability

By crafting custom URLs, any person with the ability to browse Confluence can modify content on the site, bypassing security settings. This vulnerability does not allow users to view content they would not normally be able to view, or escalate their privileges in other ways.

This flaw affects all versions of Confluence prior to 1.3.3, including the 1.4-DR development releases.

Fix

This vulnerability is fixed in Confluence 1.3.3 and later. Customers who do not wish to migrate to 1.3.3 can fix this bug using the procedure below:

1. Edit the file confluence/WEB-INF/classes/xwork.xml
2. Find the following section near the top of the file (around line 34):

```
01.<interceptor-stack name="defaultStack">
02.         <interceptor-ref name="profiling">
03.                 <param name="location">Before defaultStack</param>
04.         </interceptor-ref>
05.         <interceptor-ref name="transaction"/>
06.         <interceptor-ref name="authentication"/>
07.         <interceptor-ref name="requestParameterHack"/>
08.         <interceptor-ref name="eventnotifier"/>
09.         <interceptor-ref name="autowire"/>
10.         <interceptor-ref name="params"/>
11.         <interceptor-ref name="servlet"/>
12.         <interceptor-ref name="pageAware"/>
13.         <interceptor-ref name="permissions"/>
14.         <interceptor-ref name="profiling">
15.                 <param name="location">After defaultStack</param>
16.         </interceptor-ref>
17.</interceptor-stack>
```

3. Locate the "autowire" and "params" entries:

```
1.                               <interceptor-ref name="eventnotifier"/>
2.-->         <interceptor-ref name="autowire"/>       <--
3.-->         <interceptor-ref name="params"/>         <--
4.                               <interceptor-ref name="servlet"/>
```

4. Swap the two lines around. The whole stack should now look like this:

```
01.<interceptor-stack name="defaultStack">
02.         <interceptor-ref name="profiling">
03.                 <param name="location">Before defaultStack</param>
04.         </interceptor-ref>
05.         <interceptor-ref name="transaction"/>
06.         <interceptor-ref name="authentication"/>
07.         <interceptor-ref name="requestParameterHack"/>
08.         <interceptor-ref name="eventnotifier"/>
09.         <interceptor-ref name="params"/>
10.         <interceptor-ref name="autowire"/>
11.         <interceptor-ref name="servlet"/>
12.         <interceptor-ref name="pageAware"/>
13.         <interceptor-ref name="permissions"/>
14.         <interceptor-ref name="profiling">
15.                 <param name="location">After defaultStack</param>
16.         </interceptor-ref>
17.</interceptor-stack>
```

5. Restart Confluence.

Confluence Security Advisory 2005-12-05

A flaw has been found in Confluence by which attackers to inject malicious HTML code into Confluence. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 2.0.2

Vulnerability

By entering HTML code into the Confluence search input fields, attackers can cause arbitrary scripting code to be executed by the user's browser in the security context of the Confluence instance.

This flaw affects all versions of Confluence between 1.4-DR releases and 2.0.1.

(Atlassian was not informed of the problem before it was published by third-party security researchers. You can read the third-party security advisory here: http://secunia.com/advisories/17833/. The vulnerability was originally reported here.)

Fix

This vulnerability is fixed in Confluence 2.0.2 and later. Customers who do not wish to migrate to 2.0.2 can fix this bug using the procedure below:

1. Edit the confluence/decorators/components/searchresults.vmd
2. Replace the following reference (around line 48):

```
1.$action.getText("search.result", [$start, $end, $total, $queryString])
```

with

```
1.$action.getText("search.result", [$start, $end, $total,
$generalUtil.escapeXml($queryString)]).
```

3. Edit the confluence/search/searchsite-results.vm.
4. Replace the following reference (around line 11):

```
1.Searched for <b>$action.searchQuery.queryString</b>
```

with

```
1.Searched for <b>$generalUtil.escapeXml($action.searchQuery.queryString)</b>
```

5. Restart Confluence.

Alternatively, you can download the patched source files from CONF-4825. If you are patching a 2.0.x installation, then use the files with the .2.0 suffix. If you are patching a 1.4.x installation, then use the files with the .1.4 suffix.

Confluence Security Advisory 2006-01-20

A flaw has been found in Confluence by which attackers to inject malicious HTML code into Confluence. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 2.1.3.

Vulnerability

By entering HTML/JavaScript code into the full name of a user's profile, attackers can cause arbitrary scripting code to be executed by the user's browser in the security context of the Confluence instance.

This flaw affects all versions of Confluence between 1.4-DR releases and 2.1.2.

This issue was initally reported by Ricardo Sueiras and a fix was quickly documented by Dan Hardiker at the Confluence Community Security Advisory 2006-01-19 page. Our thanks to them for bringing this to our attention.

There is an issue in JIRA at CONF-5233.

Fix

This vulnerability is fixed in Confluence 2.1.3 and later. Customers who do not wish to migrate to 2.1.3 can fix this bug using the procedure below:

Steps to fix:

1. Copy macros.vm to your confluence/template/includes folder
2. Restart Confluence

Note: If you are using version 1.4.4, please download and copy this file instead. You will need to rename it back to `macros.vm`.

If you are not using any of the above versions, you will need to replace wrap calls to display full names of users in $generalUtil.htmlEncode(). Alternatively, send us an email. We do however encourage you to use the latest stable point release regardless of the version you are using.

Confluence Security Advisory 2006-01-23

A flaw has been found in Confluence by which the unrestricted content of a space can be revealed in search results.

Vulnerability

By entering in a space key and blank query string into the Search macro, pages from the specified space will be displayed, without filtering on page and space permissions. This can allow unpermitted users to view the excerpts of pages they don't have access to.

This flaw is confirmed to affect all releases from 1.4 to 2.1.2.

More information is available at CONF-5189.

Fix

This vulnerability is fixed in Confluence 2.1.3 and later. We strongly suggest that customers upgrade to this release to fix the vulnerability.

Customers who are using 1.4.x and do not wish to upgrade can download a patched class from CONF-5198.

Confluence Security Advisory 2006-06-14

Vulnerability

By crafting a custom HTTP request, an attacker can delete or modify global permissions settings on a Confluence site.

This flaw affects all Confluence versions between 1.4 and 2.2.2. 2.2.3 and later are not vulnerable.

Fix

This issue has been fixed in Confluence 2.2.3. Patches are also available for all versions of Confluence betwen 1.4 and 2.2.2. For more information, please see this issue report.

Atlassian STRONGLY recommends that all customers either upgrade to Confluence 2.2.3, or apply the patch.

Confluence Security Advisory 2007-07-26

In this advisory:

- Users with view permission in a space can copy and save a page
- Space name and key are not validated nor escaped

Users with view permission in a space can copy and save a page

Vulnerability

A user who has only view permissions in a space can copy a page and then save it in the space. In this way, users can create a page in a space where they have only view permission.

This flaw affects only Confluence version 2.5.4.

Fix

This issue has been fixed in Confluence 2.5.5. A patch is also available for Confluence 2.5.4. For more information, including instructions on applying the patch, please see this issue report.

If you are using Confluence 2.5.4, Atlassian **strongly** recommends that you upgrade to Confluence 2.5.5 or apply the patch.

Space name and key are not validated nor escaped

Vulnerability

The input for space name and key is not validated properly - any characters are allowed. This makes a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.5. For more information, please see this issue report.

Atlassian recommends that you upgrade to Confluence 2.5.5.

Confluence Security Advisory 2007-08-08

In this advisory:

- Input in the RSS Feed Builder is not validated
- Input when editing Space Permissions is not validated
- Number of labels that can be added to a page is not restricted
- Input when editing navigation themes is not validated
- Viewing of space content alphabetically is not validated
- Input when editing Space Name is not validated
- Input when viewing attachments by file-type is not validated

Input in the RSS Feed Builder is not validated

Vulnerability

The input for the RSS Feed Builder is not required to be escaped. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8993.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing Space Permissions is not validated

Vulnerability

The 'Grant permission to' field on the 'Edit Space Permissions' screen is not validated. This can make a Confluence instance vulnerable to an XSS or DoS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8980 and CONF-8979.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Number of labels that can be added to a page is not restricted

Vulnerability

There is no restriction on the number of labels that can be added to a page at a time. This can make a Confluence instance vulnerable to a DoS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8978.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing navigation themes is not validated

Vulnerability

The 'Navigation Page' specified in the 'Left Navigation Theme' configuration is not validated. This can make a Confluence instance vulnerable to a XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8956.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Viewing of space content alphabetically is not validated

Vulnerability

When viewing space content by alphabetic character, the input is not validated as being alphabetic. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8952.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing Space Name is not validated

Vulnerability

The 'Name' field on the 'Edit Space Details' screen is not validated. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8951.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when viewing attachments by file-type is not validated

Vulnerability

The 'Filter By Extension' field on the 'List Space Attachments' screen is not validated. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8950.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Confluence Security Advisory 2007-11-19

In this advisory:

- DWR debug mode enabled
- XSS vulnerability in exception error page
- XSS vulnerability in the URL destination for the print icon
- XSS vulnerability in wiki markup for images

Atlassian recommends that you upgrade to Confluence 2.6.1 to fix the vulnerabilities described below.

DWR debug mode enabled

Vulnerability

Debug mode was enabled by default on Direct Web Remoting (DWR). This made it easy for a potential attacker to find information about available AJAX request handlers in Confluence.

Fix

This issue has been fixed in Confluence 2.6.1. If you do not wish to upgrade at this time, you can fix the problem by editing your `<confluence install>/confluence/WEB-INF/web.xml` file. For more information, please see CONF-9718.

XSS vulnerability in exception error page

Vulnerability

The attributes and parameters were not escaped on the Confluence exception error page. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. For more information, please see CONF-9704 and CONF-9560.

XSS vulnerability in the URL destination for the print icon

Vulnerability

The print icon on the HTTP 404 error page uses the path of the requested URL, which potentially contains malicious JavaScript. The 404 page did not correctly escape it. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. A patch is supplied for customers with **Confluence version 2.6** who do not wish to upgrade at this time. For more information, please see CONF-9456.

XSS vulnerability in wiki markup for images

Vulnerability

When using image URLs in wiki markup, quotes were not correctly escaped. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. For customers with **Confluence 2.6** who do not with to upgrade at this time, the new `atlassian-renderer` JAR should resolve this issue. For more information, please see CONF-9209.

Confluence Security Advisory 2007-11-27

In this advisory:

XSS Type 2 Vulnerabilities in Macros and Wiki Markup

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed some security flaws which may affect Confluence instances in a public environment. These flaws are XSS (cross-site scripting) vulnerabilities in some of Confluence's macros and Wiki Markup, which potentially allow a malicious user (hacker) to insert their own HTML tags or script into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

Atlassian recommends that you upgrade to Confluence 2.6.2 to fix the vulnerabilities described below.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

The following macros are affected:

- {color}
- {panel}
- {section}
- {column}
- {code}

The Wiki Markup for inserting images (e.g. `!myImage.png!`) is also vulnerable to XSS exploitation.

Fix

The fix is to escape all user input, so that no user input is interpreted as HTML or CSS. In some cases we also perform stricter validation on the range of values a user can supply in an attribute.

These issues have been fixed in Confluence 2.6.2. For more information, please see CONF-9350.

> Our thanks to **Igor Minar**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2007-12-14

In this advisory:

- Vulnerability
- Fix

XSS Vulnerability in Configure RSS Feed Action

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7, or
- Download and install the patch for Confluence 2.5.8 or Confluence 2.6.2 from our JIRA site – see issue CONF-10164.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the following Confluence action:

```
1.http://www.anyhost.com/confluence/dashboard/configurerssfeed.action
```

The above Confluence action is used to build an RSS feed based on your Confluence pages and news items. The action is invoked when a selects '**Feed Builder**' from your Confluence Dashboard. It can also be invoked by simply entering the URL into the browser address bar.

Fix

These issues have been fixed in **Confluence 2.7**, which you can download from the download centre.

A patch is available for **Confluence 2.5.8** and **Confluence 2.6.2**. For more information, please see CONF-10164.

> Our thanks to **jeff peichel**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2008-01-24

In this advisory:

- XSS Vulnerability in Dashboard Action
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix

XSS Vulnerability in Dashboard Action

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.1, or
- Download and install the patch for Confluence 2.6.2 or Confluence 2.7.0 from our JIRA site – see issue CONF-10289.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the following Confluence action:

```
1.http://confluence-location/dashboard.action?spacesSelectedTab
```

The above Confluence action is used to determine which spaces are listed on a user's Dashboard. For example, the following URL requests a list of team spaces only:

```
1.http://confluence-location/dashboard.action?spacesSelectedTab=team
```

The action is invoked when a user selects one of the 'Spaces' tabs on the Dashboard, such as the '**Team**' tab. It can also be invoked by simply entering the URL into the browser address bar.

Fix

These issues have been fixed in **Confluence 2.7.1** (see the release notes), which you can download from the download centre.

A patch is available for **Confluence 2.6.2** and **Confluence 2.7.0**. For more information, please see CONF-10289.

> Our thanks to **Mary Johnson**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate her working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2008-03-06

In this advisory:

Users with View-Only Permission can Delete (Purge) Pages

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

More explanation of the ranking we chose:

- You might rank this vulnerability as **critical**, because in most installations the vulnerability will allow anonymous users to delete information.
- We have chosen a ranking of **high**, because the vulnerability does not allow privilege escalation i.e. it doesn't allow users to gain administration privileges.

Risk Assessment

We have identified and fixed a security flaw which allowed users who have 'View' permission (or higher) on a space to purge (delete) any page in that space.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.7.1** inclusive.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.2, or
- Download and install the patch for Confluence 2.6.x or Confluence 2.7.x from our JIRA site – see issue CONF-10807.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

If it is not immediately feasible to upgrade to Confluence 2.7.2 or apply a patch, we recommend an alternative strategy:

- As a temporary measure, you can block the URL which allows someone to purge (delete) a page. Please ask your website administrator to block the URL described below.
- The impact is that Space Administrators will not be able to purge individual pages or news items. However, Space Administrators can still use the 'Purge All' link to clear the entire contents of Trash.

Vulnerability

Description:
A user can use the following Confluence action to permanently delete (purge) any Confluence page, provided that the user has 'View' permission (or higher) in the space to which the page belongs:

```
1.http://confluence-location/pages/purgetrashitem.action?key=XXX&contentId=XXX
```

The above action is invoked when a space administrator clicks the 'Purge' link on the space's 'Trash' page next to a wiki page which has already been deleted.

The action can also be invoked by simply entering the URL into the browser address bar. In this way, it is possible for a user with 'View' permission (or higher) to remove a page via the 'Purge' action, even if the page has not been deleted.

Fix

These issues have been fixed in **Confluence 2.7.2** (see the release notes), which you can download from the download centre.

A patch is available for **Confluence 2.6.x**, **Confluence 2.7.0** and **Confluence 2.7.1**. For more information, please see CONF-10807.

> Our thanks to **Neeraj Jhanji**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-03-19

In this advisory:

- XSS Vulnerabilities in Various Confluence Actions
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix

XSS Vulnerabilities in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials

back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.3, or
- Download and install the patches for Confluence **2.6.x** from our JIRA site — refer to the list of issues below.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
Create, edit or copy a page or news item	From **2.2 to 2.7.2** inclusive	CONF-11027	
Add a comment	From **2.2 to 2.7.2** inclusive	CONF-11027	
Create a space	From **2.2 to 2.7.2** inclusive	CONF-11042	Wyatt Crossin
Sign up for an account	From **2.2 to 2.7.2** inclusive	CONF-11005	
Choose a page (page picker)	From **2.2 to 2.7.2** inclusive	CONF-11137	
View a user	From **2.2 to 2.7.2** inclusive	CONF-11002	
Insert an image or link	From **2.2 to 2.7.2** inclusive	CONF-11141	
Choose a user or group (user picker and group picker)	From **2.2 to 2.7.2** inclusive	CONF-11040	Jean Marois
Add a user to favourites	From **2.0 to 2.7.2** inclusive	CONF-11026	
HTTP 500 error page	From **1.3 to 2.7.2** inclusive	CONF-11019	
Add bookmark	All Confluence instances that have the Social Bookmarking plugin. Note that the plugin is bundled with Confluence since version **2.6**, so Confluence 2.6.x and 2.7.x are vulnerable even if you don't use social bookmarking. Patches are supplied for Confluence 2.6.x and 2.7.x.	CONF-11153	

Fix

These issues have been fixed in **Confluence 2.7.3** (see the release notes), which you can download from the download centre.

Patches are available for **Confluence 2.6.x**. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

> Our thanks to the people who reported some of the vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate their working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-05-21

In this advisory:

Users can Move Attachments to Any Page Regardless of Permissions

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who have 'Create Page' permission in a space to move an attachment from a page in that space to any other page in the Confluence site, regardless of the user's permissions in the destination space.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.8.0**.

Risk Mitigation

This security flaw grants extra powers only to users who already have 'Create Page' permissions in one of the spaces on the Confluence site. In most installations, this will be a trusted group of users.

If your Confluence instance allows a less trusted group of users to create and edit pages in one space, while restricting access to other spaces, you may judge it necessary to disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

Any user who has 'Create Page' permission in a Confluence space can move an attachment from a page in that space to any other page in the Confluence site, regardless of the user's permissions in the destination space.

Note: If a user has permission to create a space, they will also have 'Create Page' permission in any space they create, including a personal space. Such users could upload an attachment onto the space they have created and then move the attachment to any page in the Confluence site.

Fix

This issue has been fixed in Confluence 2.8.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.x or Confluence 2.8.0 from our JIRA site – see issue CONF-11452.

> Our thanks to **Stafford Vaughan** from CustomWare, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate it when people work with us towards identifying and solving a problem.

XSS Vulnerability in Page Information View

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.8.0** inclusive.

Risk Mitigation

If you judge it necessary, you can hide referrers on page information views by disabling this functionality.

Vulnerability

A hacker can inject their own JavaScript into the referrer URLs which are displayed on the 'Info' view of a wiki page. The rogue JavaScript will be executed when a user opens the 'Info' view.

Fix

This issue has been fixed in Confluence 2.8.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.x or Confluence 2.8.0 from our JIRA site – see issue CONF-11524.

Confluence Security Advisory 2008-07-03

In this advisory:

- XSS Vulnerability in Various Confluence Actions
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix

XSS Vulnerability in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
Create, edit or copy a page or news item	2.8.0 and 2.8.1	CONF-11985	James Rinker
Page picker and space picker	2.2.0 to 2.8.1 inclusive	CONF-11137	

Fix

These issues have been fixed in Confluence 2.8.2 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patches provided on our JIRA site. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

> Our thanks to **James Rinker** who reported some of the vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-09-08

In this advisory:

- XSS Bug: Usernames Not HTML-Encoded in All Places
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Inherited Page Restrictions Are Not Applied After 2.9 Upgrade
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in View Wiki Markup Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in Copy Page Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in Diff Page Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Bug: Usernames Not HTML-Encoded in All Places

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allowed certain users to circumvent Confluence's security measures, by including HTML markup in their own username. This could allow a malicious user to execute Javascript on another user's authenticated session.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If the user specified a username that included HTML markup (which could include Javascript), in some places Confluence would not correctly escape this source before displaying it. This could result in Javascript being executed in another user's authenticated session. To address the issue, you should update your Confluence instance as soon as possible (or follow the patch instructions on the issue).

Vulnerability

This is a classic Cross-Site Scripting issue where usernames could include malicious Javascript.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

For more information, see issue CONF-7615 which has instructions on how to patch the affected velocity template.

Inherited Page Restrictions Are Not Applied After 2.9 Upgrade

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw that caused any content permission inherited by a page to be lost during the upgrade process to Confluence 2.9.

The following Confluence versions are vulnerable: Version **2.9**; specifically instances of Confluence that were *upgraded to version 2.9* (from an earlier version) only.

Risk Mitigation

This issue can be resolved by following the steps under **Fix**, or upgrading to Confluence 2.9.1. If this cannot be done immediately, it may be prudent to manually apply restrictions to each page that is normally protected by inherited restrictions (that is, all child pages residing under a restricted page). Enacting the fix is trivial and should take around ten minutes for a typical Confluence instance.

Vulnerability

If you had given a parent page restrictions prior to the 2.9 upgrade, then any child pages that should be inheriting these restrictions are no longer restricted. This potentially renders these child pages viewable and editable by Confluence users who should not have these rights. However you should note that any space level restrictions are still respected so these affected pages are only opened as far as the space level security allows for your site. Note for individual pages where you have manually set the permissions, those pages are not at risk — just the pages underneath them using inherited permissions.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can apply the manual fix, which involves a simple series of actions in the Confluence administration screens.

For more information see issue CONF-12911.

Access Vulnerability in View Wiki Markup Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct '**View Page**' permission in a space to view the Wiki Markup source of the page content.

The following Confluence versions are vulnerable: Version **2.9** only.

Risk Mitigation

If a user knows the URL to view the source of a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string:
`/pages/viewpagesrc.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page that they do not have **'View Page**' permission for they can use the view source URL to view the Wiki Markup of a page. This will allow them to copy and paste the contents of the page to another location, or simply read the markup and deduce its final content.

Note: the user will need to know the page ID of a page. Confluence will not provide any links to the restricted page through a search or other navigation.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

For more information see issue CONF-12845.

Access Vulnerability in Copy Page Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct 'View Page' permission in a space to copy a page and therefore see its content.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If a user knows the URL to copy a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string: `/pages/copypage.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page they do not have permissions for, they can use the copy page URL to copy the page to a space where they do have permission. This will allow them to create a new page based on the content of a page they aren't meant to see.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.3 or 2.8.2 from our JIRA site – see issue CONF-12859.

Instruction on installing the patch can be found here.

Access Vulnerability in Diff Page Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct '**View Page**' permission in a space to create a diff of a page (a comparison of its contents with another page) and therefore see its content.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If a user knows the URL to perform a diff of a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string: `/pages/diffpages.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page they do not have permissions for, they can use the 'Diff Page' URL to compare the contents of that page with one where they do. This will allow them to deduce the contents of a page they don't have access to.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.3 or 2.8.2 from our JIRA site – see issue CONF-12860.

Instruction on installing the patch can be found here.

> Our thanks to **Neeraj Jhanji** from Atlassian Partner ImaHima, who reported the copy and diff page issues to Atlassian. We fully support the reporting of vulnerabilities and we appreciate it when people work with us towards identifying and solving a problem.

Confluence Security Advisory 2008-10-14

In this advisory:

- Parameter Injection Vulnerability in Confluence
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix
- XSS Vulnerability in Various Confluence Actions and Plugins
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix
- Privilege Escalation Vulnerability in Confluence Watches
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix
- Privilege Escalation Vulnerability in Confluence Favourites
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix

Parameter Injection Vulnerability in Confluence

Severity

Atlassian rates this vulnerability as **critical**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow a malicious user (hacker) to inject their own values into a Confluence request by adding parameters to the URL string. This would allow a hacker to bypass Confluence's security checks and perform actions that they are not authorised to perform.

Risk Mitigation

To address the issue, you should upgrade Confluence as soon as possible or follow the patch instructions below. If you judge it necessary, you can block all untrusted IP addresses from accessing Confluence.

Vulnerability

A hacker can design a URL string containing parameters which perform specific actions on the Confluence server, bypassing Confluence's security checks. This is because Confluence does not adequately sanitise user input before applying it as an action on the server.

Exploiting this issue could allow an attacker to access or modify data and compromise the Confluence application.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, a patch is available that will work with any affected version of Confluence. You can download and install the patch from on our JIRA site. For more information, please refer to CONF-13092.

XSS Vulnerability in Various Confluence Actions and Plugins

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
View children via the Pagetree plugin (bundled with Confluence)	2.8.0 to 2.9.1 inclusive	CONF-13043	Thomas Jaehnel
Update bookmark via the Social Bookmarking plugin (bundled with Confluence)	2.6.0 to 2.9.1 inclusive	CONF-13041	Thomas Jaehnel
Build RSS feed	2.0 to 2.9.1 inclusive	CONF-13042	Thomas Jaehnel
Search via Search macro	All versions from 1.0 to 2.9.1 inclusive	CONF-13040	Thomas Jaehnel
Search	All versions from 1.0 to 2.9.1 inclusive	CONF-12944	

Fix

These issues have been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

> Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported most of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Privilege Escalation Vulnerability in Confluence Watches

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow an unauthorised user to add a Confluence page to the list of pages they are watching, even if the user does not have permission to view that page. Under some circumstances, the unauthorised user may thus have access to information they are not authorised to see.

Risk Mitigation

This flaw does not allow the unauthorised user to update the page, but it may give the user access to information that they do not have permission to see.

Vulnerability

An unauthorised user can manipulate the HTTP request, so that it adds a watch to a page which the user does not have permission to view. The page then appears in the user's list of watched pages, displaying the page title and the corresponding space name. In this way, the user can bypass Confluence's permission checks and gain access to information they are not authorised to see.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information, please refer to CONF-13039.

> Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported the vulnerability listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Privilege Escalation Vulnerability in Confluence Favourites

Severity

Atlassian rates this vulnerability as **moderate**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow an unauthorised user to add a Confluence page to their list of favourites, even if the user does not have permission to view that page. Under some circumstances, the unauthorised user may thus have access to information they are not authorised to see.

Risk Mitigation

This flaw does not allow the unauthorised user to update the page, and it gives the user only very limited access to the information they do not have permission to see.

Vulnerability

An unauthorised user can manipulate the HTTP request, so that it marks as 'favourite' a page which the user does not have permission to view. The page is then added to the number of favourites for the user. The user cannot see the page title or content, but can see that the favourite count has been incremented.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information, please refer to CONF-13044.

> Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported the vulnerability listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2008-12-03

In this advisory:

- XSS Vulnerability in Various Confluence Actions
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability

XSS Vulnerability in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into various Confluence URLs — see the table below for the affected functional areas. A URL may be invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The URL can also be invoked by simply entering it into the browser address bar. If rogue JavaScript is injected into such a URL, the JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Affected Confluence Functionality	Affected Confluence Versions	Fix Availability	More Details	Reporter (If Not Atlassian)
Handling of error messages. (Vulnerability in the DWR code library used by Confluence.)	2.7.3 to 2.9.2 inclusive	2.9.2 and 2.10	CONF-11808	Bjoern Froebe
Attachments macro.	2.8 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10**	CONF-13713	
Uploading of attachments.	2.6 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10	CONF-13717	
Inserting images as thumbnails.	2.8 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10	CONF-13625	
Log events listed in the Confluence 500 error page.	2.9 to 2.9.2 inclusive	2.10 only	CONF-13584	
Wiki Markup link rendering.	2.7 to 2.9.2 inclusive	2.7.x, 2.8.x, 2.9.x, 2.10	CONF-13451	

* The patch for CONF-13717 also addresses the bug in CONF-13736.
** To fix this issue, please upgrade your Attachments plugin to the latest version. This plugin is available for Confluence 2.8.2, 2.9.2 and 2.10, via the Confluence Plugin Repository.

Fix

These issues have been fixed in Confluence 2.10 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.10, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.8, you will need to upgrade to version 2.8.2) and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Please note that one of the issues can **only be fixed by upgrading to Confluence 2.10**. Please see the table above for details.

> Our thanks to **Bjoern Froebe**, who reported one of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Users can View a List of All Attachments by Supplying an Edited URL

Severity

Atlassian rates this vulnerability as **medium**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows a user to view the list of all attachments for all pages in a Confluence instance, regardless of space-level or page-level permissions.

While the user cannot open the files, a range of metadata is available for viewing, including file name, the page that the file is attached to, the creator, and the creation and last-modified date of the attachment.

Risk Mitigation

If you judge it necessary, you can disable anonymous access to your wiki until you have applied the necessary patch or upgrade.

Vulnerability

If a user removes the space key from the URL while viewing attachments for a space, Confluence will display the full list of all attachments for all spaces. For more details, please refer to CONF-13874.

Fix

These issues have been fixed in Confluence 2.10 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.10, you can download and install the patches provided in the JIRA issue, CONF-13874. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.8, you will need to upgrade to version 2.8.2) and then apply the patch.

> Our thanks to **Matthew Goonan**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2009-01-07

In this advisory:

Content Overwrite Vulnerability in the Office Connector Plugin

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified a risk that makes it possible for users with read-only access to a Confluence wiki space to modify its contents via the document import feature of the Office Connector plugin. This issue, however, does not expose restricted content on a Confluence wiki space to unauthorised users.

Risk Mitigation

Please see the 'Fix' section below. If you cannot apply the fix immediately, you can consider taking one or more of the following steps:

- Disable the whole Office Connector plugin, as explained here.
- If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade.
- For even tighter control, you could restrict access to trusted groups.

Vulnerability

The Office Connector plugin was first bundled in Confluence version 2.10.0. Hence, this vulnerability affects Confluence **2.10.0** where the Office Connector Plugin is enabled. Additionally, this plugin is compatible with all versions of Confluence **from 2.3.0** onwards. Hence, if you have installed the plugin, this vulnerability will affect your Confluence instance.

Fix

Please download and install the latest version of the Office Connector plugin via the Confluence Plugin Repository (instructions here). If you wish to install this plugin manually, you can download it from here.

Alternatively, install or upgrade to Confluence version 2.10.1. (See the release notes.) The Confluence 2.10.1 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14014.

> Our thanks to **Justin Wong**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2009-02-18

In this advisory:

- HTTP Header Injection Flaw
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

HTTP Header Injection Flaw

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

ⓘ An Advanced Warning of this Security Advisory published last week on http://forums.atlassian.com, stated the severity of this vulnerability as **critical**. After further assessing the likelihood of attack, however, we have amended this to **high**.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an HTTP header injection vulnerability in the Seraph web framework that is used by Confluence. This potentially allows a malicious user (attacker) to modify the HTTP response to insert malicious code. An attacker could present a modified URL to users (e.g. disguised in an email message). If any user clicks the URL, the malicious code would be executed in the user's session.

- The attacker may take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker could also gain control over the underlying system, based on the privileges of the user whose session cookie has been stolen.
- The attacker could redirect the user to undesirable web sites. This is potentially damaging to your company's reputation.

Atlassian recommends that you upgrade to Confluence 2.10.2 to fix the vulnerabilities described below.

Risk Mitigation

We strongly recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

Alternatively, you may consider taking the following step, although the time required to fix this vulnerability and the extent of its effectiveness will depend on your application server running Confluence and its configuration:

- Consult the vendor of your application server to see whether your application server is immune to header injection vulnerabilities or has configuration options to prevent such attacks. For example, the Coyote (HTTP) connector in Tomcat version 5.5 and later is immune to header injection attacks, as acknowledged in this reference.
  Technical note: In your application server, header injection vulnerabilities can be mitigated if the setHeader(), addHeader(), and sendRedirect() methods in the HttpServletResponse class have their parameters properly checked for header termination

characters.

ⓘ You may wish to forward this technical note to the vendor of your application server to help them assess the vulnerability of your application server to header injection attacks.

Vulnerability

All versions of Confluence prior to 2.10.2 are vulnerable to this security flaw.

Fix

The fix updates the Seraph framework to a version which correctly encodes and validates redirect URLs before sending them back to the user.

To patch your existing installation of Confluence, please refer to CONF-14275. This JIRA issue contains the downloadable patch file and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.2. (See the release notes.) The Confluence 2.10.2 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14275.

Confluence Security Advisory 2009-04-15

In this advisory:

XSS Vulnerability in Various Confluence Macros

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed two security flaws which may affect Confluence instances in a public environment. These flaws are all cross-site scripting (XSS) vulnerabilities in Confluence's Index and Widget Macros. Each vulnerability potentially allows a malicious user (attacker) to embed their own JavaScript into a Confluence page, which will be executed when the page is rendered.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

Alternatively if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

You could also temporarily disable the **Widget Connector** plugin and the **Index Macro module of the Confluence Advanced Macros** plugin until you have applied the necessary patch or upgrade. Be aware, however, that this will cause any occurrence of these macros on existing pages or blogs in your Confluence site to render with 'Unknown Macro' indications.

Vulnerability

All versions of Confluence prior to 2.10.3 are vulnerable to this security flaw.

Fix

The fixes include an update to the Index Macro, such that it correctly renders content on the page and an update to the Widget Macro, such that it correctly encodes all parameters passed to it.

To patch your existing installation of Confluence, please refer to CONF-14753 for the Index Macro and CONF-14337 for the Widget Macro. These JIRA issues contain the downloadable patch files and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.3. (See the release notes.) The Confluence 2.10.3 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14753 and CONF-14337.

> Our thanks to **Igor Minar**, who reported one of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

HTTP Header Injection Flaw with Attachment Filenames

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw with attachment filenames. This vulnerability could lead to an HTTP Header Injection attack through the upload of attachments with modified filenames designed to exploit this flaw. An attacker could insert malicious code into the HTTP response, which would be executed in the user's session.

- The attacker may take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker could also gain control over the underlying system, based on the privileges of the user whose session cookie has been stolen.
- The attacker could redirect the user to undesirable web sites. This is potentially damaging to your company's reputation.

Risk Mitigation

We strongly recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

If you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Alternatively, you may consider taking the following step, although the time required to fix this vulnerability and the extent of its effectiveness will depend on your application server running Confluence and its configuration:

- Consult the vendor of your application server to see whether your application server is immune to header injection vulnerabilities or has configuration options to prevent such attacks. For example, the Coyote (HTTP) connector in Tomcat version 5.5 and later is immune to header injection attacks, as acknowledged in this reference.
  Technical note: In your application server, header injection vulnerabilities can be mitigated if the setHeader(), addHeader(), and sendRedirect() methods in the HttpServletResponse class have their parameters properly checked for header termination characters.

  You may wish to forward this technical note to the vendor of your application server to help them assess the vulnerability of your application server to header injection attacks.

Vulnerability

All versions of Confluence prior to 2.10.3 are vulnerable to this security flaw.

Fix

The fix includes a new header-injection prevention filter in Confluence, which ensures attachment filenames or any other user-provided data is correctly encoded before being included in HTTP headers.

To patch your existing installation of Confluence, please refer to CONF-14704. This JIRA issue contains the downloadable patch files and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.3. (See the release notes.) The Confluence 2.10.3 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14704.

Confluence Security Advisory 2009-06-01

In this advisory:

- XSS Vulnerability in Various Confluence Actions and Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Actions and Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. These are cross-site scripting (XSS) that affect various Confluence page/blog features and functions.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into various Confluence URLs — see the table below for the affected functional areas. A URL may be invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The URL can also be invoked by simply entering it into the browser address bar. If rogue JavaScript is injected into such a URL, the JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Affected Confluence Functionality	Affected Confluence Versions	Fix Availability	More Details
Concurrent page edit message	All versions (1.0 to 2.10.3 inclusive)	2.9.2 and 2.10.3	CONF-15883
Gallery Macro (Confluence Advanced Macros Plugin)	All versions (1.0 to 2.10.3 inclusive)	2.10.3	CONF-15376
View File Macro (Office Connector Plugin)	2.10.0 to 2.10.3 inclusive *	2.10.3	CONF-15402
Instant Messenger Macro	All versions (1.0 to 2.10.3 inclusive)	2.8.2, 2.9.2 and 2.10.3	CONF-15397
Contributors Macro	2.3 to 2.10.3 inclusive	2.9.2 and 2.10.3	CONF-15399
JIRA Issues Macro	All versions (1.0 to 2.10.3 inclusive)	2.10.3	CONF-15754

* This vulnerability may be present in earlier Confluence versions with the Office Connector plugin installed.

Fix

These issues have been fixed in Confluence 3.0 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.9, you will need to upgrade to version 2.9.2) and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Confluence Security Advisory 2009-06-16

In this advisory:

- Page Content Vulnerabilities
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Page Content Vulnerabilities

> If you have already upgraded to Confluence 3.0, then you are not affected by the vulnerabilities described on this page and there is no need to take any further action.

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed two security vulnerabilities which may affect Confluence instances in a public environment. Both of these fixes are associated with a tightening of user access restrictions when either viewing specific page content or adding new page content.

The first of these vulnerabilities allows a user without permission to view a given page, to view the contents of any files attached to that page using the view file macro. This assumes that the user has permission to edit or create another page within the Confluence site and knows the name of the file attached to the page they cannot view. For more information, please refer to the JIRA issue CONF-15809.

The second of these vulnerabilities allows users with space administrator permissions to import pages to a Confluence space. The security level of this function has been tightened to permit only users with the system administration permission to access it. For more information, please refer to CONF-15267.

Risk Mitigation

If you have not already upgraded to Confluence 3.0, then we recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

All versions of Confluence up to and including version 2.10.3 with the Office Connector plugin installed are affected by the first view file macro vulnerability.

All versions of Confluence 2.10.x are affected by the second page imports vulnerability.

Fix

These issues have been fixed in Confluence 3.0 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.10.0, you will need to upgrade to version 2.10.3) and then apply the patches. For more information, please refer to the specific JIRA issues shown below.

To download the patch to fix the first view file macro vulnerability, please refer to CONF-15809.

To download the patch to fix the second page import vulnerability, please refer to CONF-15267.

Confluence Security Advisory 2009-08-20

In this advisory:

- Privilege Escalation Vulnerability in Profile Picture Handling
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- XSS Vulnerability in Various Page and Blog Post Features and Functions

Privilege Escalation Vulnerability in Profile Picture Handling

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified a privilege escalation vulnerability, which could provide an attacker with access to administrative areas and functions of Confluence when specifying a profile picture. Under some circumstances, the attacker could gain access to Confluence administrative functions that they are not authorised to use.

Risk Mitigation

To address the issue, you should upgrade to Confluence 3.0.1 as soon as possible or follow the patch instructions in the Fix section below. If you judge it necessary, you can disable public signon to your wiki until you have applied the necessary patch or have performed the upgrade. For even tighter control, you could also restrict access to trusted groups or additionally, disable anonymous access until your system is patched or upgraded.

Vulnerability

The profile picture handling feature in all versions of Confluence up to 3.0.0 are affected by this issue. However, the Form Token Handling mechanism available in Confluence 3.0.0 and later means that the administrative areas in these versions of Confluence cannot be compromised by this vulnerability.

Fix

This issue has been fixed in Confluence 3.0.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.1 and you are running Confluence 2.10.x, you can download and install the patches provided on our JIRA site. We strongly recommend that you upgrade to the latest point release (2.10.3) before applying the patch. For more information, please refer to CONF-16141.

> Our thanks to **Elliot Kendall** of **Emory University**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

XSS Vulnerability in Various Page and Blog Post Features and Functions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of XSS vulnerabilities in various Confluence page/blog features and functions, which may affect Confluence instances in a public environment.

XSS vulnerabilities potentially allow a malicious user (attacker) to embed their own JavaScript into a Confluence page.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Clicking a username link	3.0.0	3.0.0 and 3.0.1	CONF-15970
Moving pages between spaces	2.8 to 2.10.3 inclusive	2.10.x and 3.0.1	CONF-16019 * CONF-16135 *
Entering content into the WebDAV Configuration page	3.0.0 2.10.x with version 2.0 of the WebDAV plugin	2.10.x, 3.0.0 and 3.0.1	CONF-16136
Entering content into the PDF Export Stylesheet	3.0.0	3.0.0 and 3.0.1	CONF-16209

* Applying the patch for one of these issues fixes the other.

Fix

These issues have been fixed in Confluence 3.0.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.1, you can patch your existing installation by downloading and installing the patched files provided on our JIRA site. For the WebDAV plugin vulnerability, this would involve upgrading the version of the plugin. We strongly recommend that you upgrade to the latest point release of the major version of Confluence that you are running before applying the patches. For example, if you are running Confluence 2.10.1, you should upgrade to version 2.10.3 and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Confluence Security Advisory 2009-10-06

In this advisory:

- Session Fixation Vulnerability
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- XSS Vulnerability in Various Confluence Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Session Fixation Vulnerability

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security vulnerability which may affect Confluence instances in a public environment. This vulnerability could lead to a session fixation attack, in which the malicious user (attacker) can gain access to a victim's Confluence resources whilst the victim is logged in to their Confluence user account.

The attacker does this by fixating (or setting) their session ID onto the victim's computer. While the victim is logged in, all the victim's privileges are associated with the attacker's session ID, effectively granting the attacker access to all of the Confluence data and resources accessible to the victim.

For more information about session fixation attacks, please refer to the following sources:

- Chris Shiflett's Security Corner article
- The Web Application Security Consortium's overview

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

All versions of Confluence prior to 3.0.2 are vulnerable to this security issue.

Fix

These issues have been fixed in Confluence 3.0.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.2 and you are currently running Confluence version 2.10.x or 3.0.x, you can patch your existing installation by downloading the appropriate patch file attached to JIRA issue CONF-15108 and installing the patch file using the instructions provided in this JIRA issue.

> Our thanks to **Ben L Broussard** who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

XSS Vulnerability in Various Confluence Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security vulnerabilities which may affect Confluence instances in a public environment. These flaws are cross-site scripting (XSS) vulnerabilities in Confluence's pagetree, userlister and content by label macros. These XSS vulnerabilities potentially allow an attacker to embed their own JavaScript into a Confluence page.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Pagetree Macro	2.8.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-16651
Userlister Macro	2.6.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-16644
Content by Label Macro	2.10.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-15440

Fix

These issues have been fixed in Confluence 3.0.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.2, you can patch your existing installation by upgrading the plugins for these macros via the Confluence Plugin Repository to the version indicated in the JIRA issues listed in the vulnerability section (above).

Confluence Security Advisory 2009-12-08

In this advisory:

- XSS Vulnerability in Various Confluence Actions and Macros
    - Severity
    - Risk Assessment
    - Risk Mitigation
    - Vulnerability
    - Fix

XSS Vulnerability in Various Confluence Actions and Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Severity Levels for Security Issues. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security vulnerabilities which may affect Confluence instances in a public environment. These flaws are cross-site scripting (XSS) vulnerabilities that could occur when creating a page or blog post in a personal space, using the `indexbrowser.jsp` form and when using the gallery macro.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Page or blog post creation in a personal space	2.10 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17031
Using the `indexbrowser.jsp` form	All versions prior to and including 3.0.2	3.0.0 – 3.1 inclusive	CONF-17165
Gallery macro	2.9 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17361
Page tree and page tree search macros	2.9 – 3.0.2	2.8 – 3.1 inclusive	CONF-17967
Status updates tab of the user profile area	3.0.0 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17933

Fix

These issues have been fixed in Confluence 3.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.1, you can patch your existing installation by upgrading the plugins for these macros via the Confluence Plugin Repository to the version indicated in the JIRA issues listed in the vulnerability section (above).

Design and Layout

RELATED TOPICS

Modifying Confluence Interface Text
Site Configuration

Custom Decorator Templates

About Decorators

Confluence is built on top of the Open Source SiteMesh library, a web-page layout system that provides a consistent look and feel across a site. SiteMesh works through "decorators" that define a page's layout and structure, and into which the specific content of the page is placed. If you are interested, you can read more on the SiteMesh website.

What this means for Confluence is that you can customise the look and feel of almost all of your Confluence site through editing three decorators:

- The "Main" decorator defines the look and feel of most pages on the site
- The "Popup" decorator defines the look and feel of the popup windows such as the "Insert Link" and "History" pages.

- The "Printable" decorator defines the look and feel of the printable versions of pages (available through the 🖨 icon on each page)

You can view and edit these decorators from within Confluence: they are available from the "Layouts" option on the site's Administration menu. Changes to the decorators will affect all spaces hosted on that Confluence installation.

The decorator that is used to draw Confluence's administrative pages can not be edited from within Confluence. This means that if you make some editing mistake that renders the rest of the site unuseable, the administrative pages should still be available for you to fix the template.

Browsing the Default Decorators

At any time, you can browse the default decorators that come packaged with Confluence by following the "View Default" links on the "Site Layouts" page. The template browser also allows you to view the "#parsed" templates that are included within the template when it is compiled. While you can't edit these included templates, you will probably have to copy some or all of them into your custom template as you do your customisation.

Editing Custom Decorators: Add a Logo

To edit Confluence decorators, you should have a good knowledge of HTML, and some understanding of the Velocity templating language.

The first thing you will see when you choose to create a custom "Main" decorator is... there's not much to edit. By default, most of the content

of this decorator is included from other files:

```
01.<html>
02.<head>
03.            <title>$title - Confluence</title>
04.
05.            #standardHeader()
06.</head>
07.<body onload="placeFocus()">
08.
09.<div id="Content">
10.            <table border="0" cellpadding="0" cellspacing="0" width="100%">
11.                    <tr>
12.                            <td width="60%" rowspan=2 class="logocell"
>#pagetitle("spacenametitle")</td>
13.
14.                            <td width="40%" align="right" valign="top"
>#globalnavbar("table")</td>
15.                    </tr>
16.                    #if ($setup.isSetupComplete())
17.                    <tr>
18.                            <td align=right valign="bottom">
19.                                    #usernavbar()
20.                                    #printableicon()
21.                                    #helpicon()
22.                            </td>
23.                    </tr>
24.                    #end
25.            </table>
26.
27.            #breadcrumbsAndSearch()
28.
29.            <table border="0" cellpadding="5" cellspacing="0" width="100%"><tr><td>
30.<table border="0" cellpadding="0" cellspacing="0" width="100%"><tr>
31.
32.            <td valign="top" class="pagebody">
33.## The "toolbar-style" page operations
34.## #if ($page.getProperty("page.operations"))
35.## <table align="right" class="toolbar"><tr><td>
36.## $page.getProperty("page.operations")
37.## </td></tr></table>
38.## #end
39.
40.                    #if ($page.getProperty("page.surtitle"))
41.                            $page.getProperty("page.surtitle")
42.                    #end
43.
44.                    #if (!$page.getProperty("page.no-page-header"))
45.                            <div class="pageheader">
46.                                    <span class="pagetitle">$title</span>
47.                            </div>
48.                    #end
49.
50.                    $body
51.            </td>
52.
53.#parse ("/decorators/includes/complete_footer.vmd")
```

We can add our logo, changing the "logocell" table cell:

```
1.<td width="60%" rowspan=2 class="logocell">
2.<img align="right"
3.src=http://www.atlassian.com/images/atlassian_logo.gif
4.width="203" height="60">#pagetitle("spacenametitle")</td>
```

When you insert this into the right section of the template and hit save, visitors to the site will see the logo at the top of each page. Note, the administrative pages will be unaffected: you will have to go to the dashboard or to a space to see the changes you have made.

Macros

Some parts of the page are drawn using Velocity macros, including the navigation bar. The macros you should know about when editing decorators are described in Working With Decorator Macros.

If Something Goes Terribly Wrong

From the "Site Layouts" page in Confluence's administrative menu, you can delete your custom templates. When you do this, the default template will be restored, fixing anything that may have been broken.

Alternatively, the custom templates are stored inside the confluence.home directory you defined in confluence-init.properties when installing Confluence. If you have somehow managed to render Confluence completely unuseable through editing your templates,

simply delete the `confluence.home/velocity` directory, and restart Confluence. The default templates will be restored.

WARNING: Only delete the `velocity` directory! Changing anything else inside your `confluence.home` is dangerous, and you could lose important data!

For Advanced Users

The `velocity` directory is at the front of Confluence's velocity template search path. As such, you can override *any* of Confluence's velocity templates by placing an identically named file in the right place.

While we don't recommend you do this unless you know exactly what you're doing, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through the web interface.

There are, however, two important caveats:

1. Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.
2. Because we only officially support the modification of the three global decorator files, other changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site.

Customising Look and Feel Overview

You can customise the 'look and feel' of Confluence at both the global and space levels.

Any changes you make to the look and feel of the site at the global level will be applied as the default look and feel for all the spaces in the site. This means that any customisations will only be reflected in the "Default" theme. No other theme will have an impact from this change. An individual space can be configured to have its own look and feel through the space administration screens.

Here's how you can customise the look and feel of your site:

- **Colour Scheme** : Change the colour scheme of the user interface.

- **Layouts** : Edit how the controls are laid out in the site. This does not change the actual page layouts but the way the surrounding controls appear in the page.

- **Themes** : Use themes for advanced layout customisation.

RELATED TOPICS

📄 Adding a Site-Wide Banner

📄 Customising Look and Feel Overview

📄 Editing the Footer

📄 Customising Layouts

📄 Global Templates

📄 Upgrading Custom Layouts

📄 Customising Colour Schemes

📄 Working With Decorator Macros

📄 Customising a Specific Page

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Customising Colour Schemes

A Confluence administrator can configure a new colour scheme for the site dynamically from the **Administration Console**.

The default colour scheme for the site will also become the default for all spaces within it. However, it is possible for space administrators to configure a different colour scheme for spaces from the space administration screens.

To change the site's colour scheme,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Colour Scheme**' in the left-hand panel.
   This will bring up a new screen. See screenshot below.

3. Click '**Edit**'. Enter standard HTML/CSS2 colour codes, or use the colour-picker to choose a new colour from the palette provided. Any changes you make will immediately be reflected across the Confluence installation.

The colour scheme applies to the following UI elements:

- **Top Bar** - the bar across the top of the page that contains the breadcrumbs
- **Tab Navigation Background** - the background colour of the tab navigation menus
- **Tab Navigation Text** - the text of the tab navigation menus
- **Breadcrumbs Text** - the breadcrumbs text in the top bar of the page
- **Space Name Text** - the text of the current space name located above the page title
- **Heading Text** - all heading tags throughout the space.
- **Links** - all links throughout the space.
- **Borders and Dividers** - table borders and dividing lines.
- **Tab Navigation Background Highlight** - the background colour of the tab navigation menu when highlighted
- **Tab Navigation Text Highlight** - the text of the tab navigation menu when highlighted
- **Top Bar Menu Selected Background** - the background colour of the top bar drop down menu when selected
- **Top Bar Menu Item** - the text colour of the menu items in the top bar drop down menu
- **Page Menu Selected Background** - the background colour of the drop down page menu when selected
- **Page Menu Item Text** - the text of the menu items in the drop down page menu
- **Menu Item Selected Background** - the background colour of the menu item when selected (applies to both the top bar and page drop down menus)
- **Menu Item Selected Text** - the text colour of the menu item when selected (applies to both the top bar and page drop down menus)

Please note that some UI elements are specific to the default theme and may not take affect for other themes.

Screenshot : Editing a site's colour scheme

Custom Colour Scheme
A custom colour scheme which can be edited.

Selected

Edit

The following colours can be customised for this colour scheme.

Top Bar	#003366
Tab Navigation Background	#3c78b5
Tab Navigation Text	#ffffff
Breadcrumbs Text	#ffffff
Space Name Text	#999999
Heading Text	#003366
Links	#003366
Borders and Dividers	#3c78b5
Tab Navigation Background Highlight	#003366
Tab Navigation Text Highlight	#ffffff
Top Bar Menu Selected Background	#336699
Top Bar Menu Item Text	#003366
Page Menu Selected Background	#6699cc
Page Menu Item Text	#535353
Menu Item Selected Background	#6699cc
Menu Item Selected Text	#ffffff

Reset | Save | Cancel

> ✅ **Handy Hint**
> If you mess things up, just click the '**Reset**' button and then try again.

RELATED TOPICS

📄 Adding a Site-Wide Banner

📄 Customising Look and Feel Overview

📄 Editing the Footer

📄 Customising Layouts

📄 Global Templates

📄 Upgrading Custom Layouts

📄 Customising Colour Schemes

📄 Working With Decorator Macros

📄 Customising a Specific Page

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Customising Layouts

> ⛔ If you modify the look and feel of Confluence by following these instructions, you will need to update your customisations when upgrading Confluence. The more dramatic the customisations are, the harder it will be to reapply your changes when upgrading. Please take this into account before proceeding with your customisation.

You can customise the layout of your Confluence instance by editing the 'decorators' that define the look and feel of the site. You need to have System Administrator permissions in order to perform this function.

Confluence is built on top of the open source SiteMesh library, a web-page layout system. Read more on the SiteMesh website. To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a .vmd file and is written in a very simple programming language called Velocity. You can learn more from the Velocity User Guide.

Once you are familiar with Velocity, you can edit the decorator files to personalise the appearance of Confluence.

> ℹ️
> - You can customise the layouts for a particular space or for the whole site. This page tells you how to customise layouts for the site as a whole. To customise the layouts for a space, use the 'Layout' menu on the 'Space Admin' page.
> - When you upgrade Confluence, you must reapply your custom layouts to the newly installed default layouts.

The decorator files are grouped into:

- **Site layouts** : These are used to define the controls that surround each page in the site. For example, the header and the footer.

- **Content layouts** : These control the appearance of content such as pages and news items: they don't change the way the pages themselves are displayed, but allow you to alter the way the surrounding comments or attachments are displayed.

- **Export Layouts**: These control the appearance of spaces and pages when they are exported to HTML. If you are using Confluence to generate a static website, for example, you will need to modify these layouts.

Editing a site decorator file

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Layouts**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
   - Click '**View Default**' to view the vmd file.
   - Click '**Create Custom**' to edit the default vmd file. This will open up the vmd file in edit mode.
3. Make changes and click '**Update**'.

ⓘ **If something goes wrong** : Click '**Reset Default**' to revert to the original layouts.

Using Velocity macros

When editing Custom Decorator Templates, there are a number of macros available to define complex or variable parts of the page such as menus and breadcrumbs. You may insert these macros anywhere in your templates. More information on Working With Decorator Macros.

For advanced users

The `velocity` directory is at the front of Confluence's velocity template search path. As such, you can override *any* of Confluence's velocity templates by placing an identically named file in the right place. While we don't recommend you do this unless you know exactly what you're doing, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through the web interface.

> ⚠ **Caching**
> Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

In Confluence 2.6 and later, some Velocity files are located inside the Confluence JAR file that can be found at `confluence/WEB-INF/lib/confluence-x.x.x.jar`. To override files inside this JAR (which you can open with any ZIP tool like WinZip or 7-Zip), put your customised file in the same directory structure under `confluence/WEB-INF/classes/`.

For example, the file `templates/macros/alphaindex.vm` inside confluence.jar can be replace by putting your custom file in `WEB-INF/classes/templates/macros/alphaindex.vm`. You do not need to modify the file inside the JAR.

See also Editing Files within JAR Archives.

RELATED TOPICS

📄 Adding a Site-Wide Banner

📄 Customising Look and Feel Overview

📄 Editing the Footer

📄 Customising Layouts

📄 Global Templates

📄 Upgrading Custom Layouts

📄 Customising Colour Schemes

📄 Working With Decorator Macros

📄 Customising a Specific Page

Velocity Template Overview
Basic Introduction to Velocity

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Adding a Navigation Sidebar

> ⚠ **This customisation must be re-applied each time you upgrade Confluence**
> When you upgrade to a new major Confluence version (e.g. from Confluence 2.9.x to Confluence 2.10.x or from Confluence 3.0.x to Confluence 3.1.x), you will need to re-apply this customisation. See instructions below.

Confluence pages can be set to include a left-hand navigation sidebar (table of contents) as shown below:

You can see an online example on the left of this page, and in other documentation spaces such as the Crowd documentation.

> **ⓘ Information for people using Confluence 2.8 or earlier**
> The instructions below will work for Confluence 2.9 and later only. For earlier versions, please see the navigation panel on the left of this page or go to the documentation home page, select the documentation space for your version under 'All Versions' and then find the page called 'Adding a Navigation Sidebar' within that space.

Follow the instructions below to add the navigation sidebar to your Confluence space.

> **⚠ Permissions required**
> To customise a space layout as described below, you must be a Space Administrator in the given space and you must be a System Administrator on the Confluence site.

Step 1. Create the TreeNavigation Page

First, you will create a Confluence page containing the `pagetree` macro. This is just a normal Confluence page. The only slight oddity is that it should reside at the root of your space, instead of under the space's home page.

Follow these instructions:

1. Go to the '**Space Pages**' view for the current space. To do this:

   - Go to a page in the space, open the '**Browse**' menu and select '**Pages**'. The 'Space Pages' view will open.
   You are now at the 'root' level of your space. The 'root' level contains pages that are added above the space's home page, not as children of the home page.
2. At the root level of the space, create a page named 'TreeNavigation'.
3. On the page, insert the following text:

   ```
   1.{pagetree}
   ```

4. Now decide if you want to add extra functionality to your page tree. By default, using the code above, the page tree will use the home page of the space as its root. You can choose to:
   - Specify a different root for your page tree.
   - Add a search box at the top of the tree.
   - Allow the viewers to expand and collapse the whole tree.
   - Control other aspects of the display.
     For more information, read about the Pagetree macro.

Step 2. Change the Page Layout on your Space

Now you will change the page layout on your space, to include the above page on the left of every web page displayed.

1. Go to the '**Space Admin**' tab of the Browse Space view. To do this:

   - Go to a page in the space, open the '**Browse**' menu and select '**Space Admin**'. The 'Space Administration' view will open.
     ℹ '**Space Admin**' is only displayed if you are a space administrator.

2. Make sure the Confluence Default theme is selected from the '**Themes**' menu.

3. Click '**Layout**' under the 'Look and Feel' section.
   ℹ 'Layout' is only displayed if you are a Confluence Site Administrator.

4. Click '**Create Custom**' under the 'Page Layout' section.

5. In the layout, locate the 'VIEW' section, and find this code:

```
1.<div class="wiki-content">
2.$body
3.</div>
```

6. Replace the above code block with this code:

```
01.#if ($action.isPrintableVersion() == false)
02.<style>
03..spacetree * ul{
04.padding-left:0px;
05.margin-left: 0px;
06.}
07..spacetree * li{
08.margin-left: 5px;
09.padding-left:5px;
10.}
11.
12.</style>
13.
14.<table cellspacing="2" cellpadding="5">
15.<tr>
16.<td valign="top" align="left" width="22%" bgcolor="#F9F9F9" class="noprint">
17.<div class="tabletitle">Table of Contents</div>
18.<div class="spacetree">
19.#includePage($helper.spaceKey "TreeNavigation")
20.</div>
21.</td>
22.<td valign="top" align="left" width="78%" class="pagecontent">
23.<div class="wiki-content">
24.$body
25.</div>
26.</td>
27.</tr>
28.</table>
29.#else
30.<div class="wiki-content">
31.         $body
32.</div>
33.#end
```

7. If you want to, you can change the table title in the above code from 'Table of Contents' to something else. For example, it might say 'Confluence Documentation'.

8. Save the updated layout.

Re-Applying the Customisation on Upgrade

When you upgrade to a new major Confluence version (e.g. from Confluence 2.9.x to Confluence 2.10.x or from Confluence 3.0.x to Confluence 3.1.x), you will need to re-apply this customisation.

Reason:
The new Confluence version may contain updates to the page layouts. Because you have customised the page layouts, Confluence will not overwrite your customisation. So your space will not get the latest updates until you set the layout to default and then re-apply your changes.

Here's how to do it:

1. First make a copy of your customised code, if you have changed it from the code above:
   - Go to '**Space Admin**', click '**Layout**' and edit the customised page layout (as created above).
   - Copy the section of code that inserts the customised left-hand navigation panel.
   - Close the page layout.
2. Click '**Reset Default**' next to '**Page Layout**', to set the page layout back to default. This will bring in the new code for the upgraded version of Confluence.
3. Create a custom page layout as described in step 2 above, and reinsert the custom left-hand navigation code.
4. Save the updated layout.

> ℹ️ **The 'All Versions' section in the navigation bar**
> A number of people have asked how we created the '**All Versions**' section at the top of our navigation side bar. Take a look at Adding an All Versions Section to your Navigation Bar.

RELATED TOPICS

Customising Layouts
Upgrading Custom Layouts
Customising the Left Navigation Theme
Example Confluence Designs

Adding an All Versions Section to your Navigation Bar

This page gives an example of how you might add an '**All Versions**' section to your navigation side bar, as currently used in the Confluence documentation, Crowd documentation and the other Atlassian product documentation spaces.

If you are viewing this page online on the Atlassian documentation wiki, you will be able to see the '**All Versions**' section at the top left of the navigation sidebar. Below is a screenshot.

A number of people have asked how we do it, so this page gives the answer. For details about creating the navigation side bar itself, please refer to Adding a Navigation Sidebar.

> ✅ **Hint: Viewing the Source Code of a Page**
> To see the Wiki Markup for one of the Atlassian documentation pages, open the '**Tools**' menu and select '**View Wiki Markup**'. You will see the macros and other markup used to create a page.

Screenshot: 'All Versions' section (expanded) at top left of navigation bar

Adding the Version Index to the Navigation Sidebar

This is how we added the '**All Versions**' section to the sidebar:

- For each product (Confluence, Crowd, Bamboo, etc) there is a page in the Inclusions Library of the ALLDOC space. The page lists all the versions of that product's documentation, linking to the relevant spaces. For example, here is the page for Confluence and the page for Crowd.

    We put the 'all versions' page in ALLDOC because the page is used in a number of different spaces, via the {include} macro. For example, the 'all versions' page may be included:
    - In every documentation space (each version) for the product concerned, such as DOC, CONF29, CONF28, CROWD, CROWD013, CROWD012, etc.
    - In the Enterprise Hosting doc space.
    - As a panel on the documentation home page, as shown in the 'All Versions' panel of the above screenshot, as well as in the left-hand navigation bar.
    - Any other places where useful.
- In each documentation space, there is a page called 'TreeNavigationVersions' like this one or this one, which copies in the content of the above 'all versions' page.
- For each documentation space, the space's page layout now includes two pages instead of just one:
    - The 'TreeNavigation' page, as already described on the page above.
    - The new 'TreeNavigationVersions' page.

Here's the relevant section of our page layout as it is currently for the Confluence documentation (DOC) space:

```
01.#if ($action.isPrintableVersion() == false)
02.<style>
03..spacetree * ul{
04.padding-left:0px;
05.margin-left: 0px;
06.}
07..spacetree * li{
08.margin-left: 5px;
09.padding-left:5px;
10.}
11.
12.</style>
13.
14.<table cellspacing="2" cellpadding="5">
15.<tr>
16.<td valign="top" align="left" width="30%" bgcolor="#eeecec" class="noprint">
17.<div class="tabletitle">All Versions</div>
18.<div class="spacetree">
19.#includePage($helper.spaceKey "TreeNavigationVersions")
20.</div>
21.<div class="tabletitle">Confluence 2.10 Documentation</div>
22.<div class="spacetree">
23.#includePage($helper.spaceKey "TreeNavigation")
24.</div>
25.</td>
26.<td valign="top" align="left" width="70%" class="pagecontent">
27.<div class="wiki-content">
28.$body
29.</div>
30.</td>
31.</tr>
32.</table>
33.#else
34.<div class="wiki-content">
35.          $body
36.</div>
37.#end
```

Adding the Expand/Collapse Functionality to the Version Index

Another question we are asked is how we group the content of the included page under a collapsible control or 'twisty'.

We use the {expand} macro. This is a 'user macro', which means that you can add it to your Confluence site by adding the code into the '**User Macros**' section of your Confluence Administration Console. The details are on the Expand macro's documentation page.

RELATED TOPICS

Adding a Navigation Sidebar

Upgrading Custom Layouts

As Confluence evolves, so do the default layouts that drive the rendering of every page. As new functionality is added or current functionally is changed, the default layouts are modified to support these changes.

> ⚠ If you are using custom layouts based on defaults from a previous Confluence version, you run the risk of **breaking functionality**, or worse, **missing out on great new features**!

Take care on each new release of Confluence to reapply your changes to the new default templates.

To reapply your custom layouts, you need to:

1. Obtain the source of your custom layouts from your current version of Confluence.
2. Reapply your customisations to the new default layouts.

Step 1. Obtaining your Custom Layouts

Ideally, you should keep a record of each customisation you have applied to each of your Confluence site or space layouts.

If not, you should be able to find your customisations using the following method.

Before Confluence 2.3, custom layouts are stored in the `velocity` directory within your Confluence home directory tree. You can open these files in any text editor.

In Confluence 2.3 and later, custom layouts are stored in the `DECORATOR` table within your Confluence database. You can `SELECT` for the source of the layout using SQL like this:

```
mysql> select SPACEKEY,DECORATORNAME,BODY from DECORATOR;
+----------+--------------------+------+
| SPACEKEY | DECORATORNAME      | BODY |
+----------+--------------------+------+
| NULL     | decorators/main.vmd | ...  |
+----------+--------------------+------+
1 row in set (0.03 sec)
```

This example was tested on **MySQL**, but should be applicable to all SQL databases.

Step 2. Reapplying your Customisations

When you upgrade Confluence to another major release of Confluence, you will need to manually re-apply any customisations you made to any site-wide or space-specific layouts. Unless otherwise stated, you should not need to re-apply customisations after conducting a minor release upgrade of Confluence.

> **What are 'major' and 'minor' release upgrades?**
>
> Major release upgrades are ones where the 1st digit of Confluence's version number or the 1st digit after the 1st decimal place differ after the upgrade, for example, when upgrading from Confluence 3.0 to 3.1, or 2.8 to 3.0. Minor release upgrades are ones where the 1st digit of Confluence's version number and the 1st digit after the 1st decimal place remain the same after the upgrade, for example, when upgrading Confluence 3.0 to 3.0.1.

If you have made Confluence site-wide layout customisations:

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Layouts**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click '**Reset Default**' next to the layout whose customisations need to be reapplied.
5. Click '**Create Custom**' next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the '**Save**' button.
7. Repeat this procedure from step 4 for each layout whose customisations need to be reapplied.

If you have made space-specific layout customisations:

1. Visit any page in the relevant space.

2. Go to the '**Space Admin**' tab of the Browse Space view. To do this:

    - Go to a page in the space, open the '**Browse**' menu and select '**Space Admin**'. The 'Space Administration' view will open.

        'Space Admin' is only displayed if you are a space administrator.
3. Select '**Layout**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
4. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
5. Click '**Reset Default**' next to the layout whose customisations need to be reapplied.
6. Click '**Create Custom**' next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
7. Click the '**Save**' button.
8. Repeat this procedure from step 5 for each layout whose customisations need to be reapplied.

Turning off caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

For Confluence 2.6, the `velocity.properties` file is available in the `confluence-2.6.0.jar` file. The jar file is located in the `WEB-INF/lib` directory. If you wish to make modification to the files in the jar, we recommend the following steps:

1. Stop Confluence.
2. Make a backup copy of the jar file.
3. Un-jar the file
4. Locate and edit the appropriate file that you wish to modify.
5. Re-jar the confluence-2.6.0.jar file.
6. Relocate the jar file to the appropriate directory.
7. Restart Confluence.

> ⛔ **Test your modifications carefully**
> Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site. It's beyond the scope of Atlassian Support to test and deploy these changes.

Global Templates

A template is a pre-defined page that can be used as a prototype when creating new pages. Templates are useful for giving pages a common style or format.

Global Templates are defined by Confluence administrators and are available in every space across the site.

Templates are written in regular Confluence markup, using special markup to define form fields that need to be filled in.

To add a global template,

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global templates**' in the left navigation panel.
3. Click '**Add new global template**'.
4. Enter a name for your template in the '**Name**' text field and an optional description in the '**Description**' text field.
5. Using regular Confluence markup and form field markup (if you are using forms), enter content in the text-entry box as you would in any other Confluence page.
6. Click '**edit**' next to '**Labels**' if you want to use labels to categorise information. Add your labels. These labels will be included in all pages created using this template.
7. Preview and click '**Save**'.

Screenshot: A template as used to create a page

RELATED TOPICS

📄 Creating a Page using a Template

📄 Form Field Markup for Templates

Modify Confluence Interface Text

All Confluence UI text is contained in a single Java properties file. This file can be modified to change the default text, and also to translate Confluence into other languages than English.

The UI text file is `ConfluenceActionSupport.properties`. From your Confluence install directory:

```
\confluence\WEB-INF\lib\confluence-3.x.jar

Within this File, the relevant file to edit is
:\com\atlassian\confluence\core\ConfluenceActionSupport.properties.
```

Refer to Editing jar files for reference.

The file contains parameters with `name=value` pairs, in the format:

```
parameter.name=Parameter value
```

Parameter names are any text before the '=' character and should never be modified. Any text after the '=' character is the parameter value, which can be modified freely and can also contain variables. An example involving variables is:

```
popular.labels=The three most popular labels are {0}, {1} and {2}.
```

For more information on replacing values, check out Translating ConfluenceActionSupport Content. Note that plugins store their text internally, so you must modify plugin text individually.

Steps For Modification

1. Stop Confluence
2. Under your install directory, open
   `\confluence\WEB-INF\lib\confluence-3.x.jar\com\atlassian\confluence\core\ConfluenceActionSupport.p`
3. Search for the text you wish to modify, replace it and save the file in
   `<Confluence-Install>\confluence\WEB-INF\classes\com\atlassian\confluence\core`. Please create this folder structure, if it does not exist already.
4. Restart Confluence

Common Modifications

- Rename 'Dashboard' by searching for `Dashboard`. To change "Dashboard" to "My Portal", change `dashboard.name=Dashboard` to `dashboard.name=My Portal`

Common Modifications

Task	Search For	Notes
Rename 'Dashboard'	Dashboard	The `dashboard.name` parameter has the name. To change 'Dashboard' to 'My Portal', change `dashboard.name=Dashboard` to `dashboard.name=My Portal` and update any other occurrences of the word 'Dashboard' in the instance
Modify login page text	login.	The `login.instructions` parameter has the "Enter your account details below to login to Confluence" text

Modify Keyboard Shortcuts

Confluence provides a set of keyboard shortcuts. You could customise the shortcuts by making modifications inside the `ConfluenceActionSupport.properties` file.

- To disable a particular shortcut, you can simply just comment out a respective line of code. One may like to disable the shortcut to one of the navigation links: *View, Edit, Attachments, Info* . For instance, to disable shortcut to *Attachments* one would comment out the following line:

```
1.#navlink.attachments.accesskey=a
```

- To modify an access key, one could simply just change the letter, bearing in mind the fact that the letter must be unique.

Working With Decorator Macros

Decorator Macros are Velocity macros which are used to draw complex or variable parts of the page such as menus and breadcrumbs when editing Custom decorators. Decorator macros can be inserted anywhere in your templates.

The macro is called by inserting a string of the form: #macroName("argument1" "argument2" "argument3").There are no commas between the arguments. Unless otherwise noted, these macros take no arguments.

NOTE: These macros will only work reliably when customising `main.vmd`. They may not work in other Velocity decorators. Decorator macros will not work inside normal confluence pages.

Macro	Usage
`#breadcrumbs()`	Draws the "You are here" breadcrumbs list, like the one found above the page name in the default template.
`#includePage(pageTitle)`	Includes a confluence page with the specified title. If you have 2 or more pages with the same title across multiple spaces, this macro will include the page belonging to the space you are currently viewing.
`#searchbox()`	Inserts a search box into the page, like the one to the far right of the breadcrumbs in the default template.
`#globalnavbar(type)`	Draws the global navigation bar, as found in the top right-hand corner of the default template. The navigation bar can be displayed in two modes:
`#globalnavbar("table")`	Displays the navigation bar in its default mode: drawn as a table of links with coloured backgrounds and mouse-over effects.
`#globalnavbar("text")`	Displays the navigation bar as series of text links separated by | characters.
`#usernavbar()`	Draws the user-specific navigation-bar. This bar contains the links to the user's profile and history, or to the login and signup pages if the user is not logged in.
`#helpicon()`	Draws the help icon, and link to the Confluence help page.
`#printableicon()`	On pages where a printable version is available, draws the printable page icon, linking to the printable version of the page. Otherwise, draws nothing
`#pagetitle(class)`	When you are viewing a page in a Confluence space, draws the name of the space that page is in. Otherwise, writes the word "CONFLUENCE".The "class" argument is the CSS class that the title should be drawn in. Unless you have customised your Confluence installation's CSS file, you should call this with "spacenametitle" as the class: `#pagetitle("spacenametitle")`
`#poweredby()`	Writes out the "Powered by Confluence" and Confluence version-number boilerplate found at the bottom of the default template.
`#bottomshadow()`	Draws the fading shadow-effect found at the bottom of the content area in the default template.
`#dashboardlink()`	Inserts a link to the dashboard page.

RELATED TOPICS

Enabling HTML macros

Include Page Macro

Editing and Removing macros

Enabling the html-include Macro

Customise pdf or html content

To customise the pdf output, modify confluence-x.y.z-jar/com/atlassian/confluence/pages/Page.pdfexport.vm. See Editing Files within JAR Archives to learn how to repackage this file.

For html output, the file is Page.htmlexport.vm in the same directory.

Check CONF-5519 for how to remove the creator's name.

Customising a Specific Page

If you'd like to change the appearance of a specific page, you can modify the corresponding Velocity template. Here's how to find out which one:

1. Access the page. Note the name of the action. For example, the "Contact Administrators" page is `<baseUrl>/administrators.action`.
2. Browse to <confluence-install>/confluence/WEB-INF/lib/confluence-x.y.jar. Copy the file.
3. Unzip or unjar the file using a standard unzipper or the java jar utility.
4. Open xwork.xml. Search the file for the name of the action corresponding to the page you'd like to modify. You'll see an entry like:

```
1.<action name="administrators" class=
"com.atlassian.confluence.user.actions.AdministratorsAction">
2.            <interceptor-ref name="defaultStack"/>
3.            <result name="success" type="velocity">/administrators.vm</result>
4.</action>
```

5. The file to look for is the vm or vmd file. In the above example, it's administrators.vmd. Because there is no context path (just a / before the name of the file), its in the root of the Confluence webapp. For the stand-alone, that's <confluence-install>/confluence folder.
6. Modify the file.

For details on how to configure the file, check the Velocity Template Overview.

RELATED CONTENT

 Adding a Site-Wide Banner

 Customising Look and Feel Overview

 Editing the Footer

 Customising Layouts

 Global Templates

 Upgrading Custom Layouts

 Customising Colour Schemes

 Working With Decorator Macros

 Customising a Specific Page

Customising the dashboard for Administrators

> To make modifications to the dashboard, modify the global template `/confluence/decorators/global.vmd` or the layout at Administration >> Layouts >> Global Layout.

For example, search the Global Layout for these macros:

```
$helper.renderConfluenceMacro("{recently-updated-dashboard:dashboard|showProfilePic=true|types=page,b
$helper.renderConfluenceMacro("{favpages:maxResults=$maxFavouritePages}")
```

To modify the bundled plugin confluence dashboard macros:

1. Modify the atlassian-bundled-plugins.zip file located at <Confluence install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup.
2. Update confluence-dashboard-macros-x.x.jar file, rezip it and then put it back to <Confluence install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup. Refer to Editing Files within JAR Archives.
3. Wipe the jar from <confluence-home>/bundled-plugins and restart.

To customise the space list, you can work with spacelist.vm.

RELATED TOPICS

Customising Look and Feel Overview

Customising the eMail Templates

> Customisations to the Confluence email templates will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.

Only administrators with access to the server where Confluence is running can modify the Confluence email templates.

Process to change the email templates

1. Shut down your test instance of Confluence.
2. In the Confluence web application folder, find the file `/confluence/WEB-INF/lib/confluence-2.x.jar`.
3. Make a copy of this file as a backup.
4. Learn how to edit files within .jar archives.
5. Within the jar file, find the `/templates/email` folder. Find the appropriate file(s) within that folder.
6. Edit the file with a text editor to make the required changes. The content is mostly HTML, but has some Velocity template variables in it. See Velocity Template Overview for more information about how these work.
7. Again using the guide on editing files within .jar archives, either rejar the set of folders or drop the new files into the identical folder structure in the `WEB-INF/classes` directory.
8. Start Confluence up again and test your changes.
9. Apply the changes to your production Confluence instance.

The same process can be applied to modify most of the templates in the Confluence web application. For velocity files that are not in a jar file, you need not shut down and restart Confluence. Be careful to test your changes before applying them to a live site. The templates contain code that is vital for Confluence to function, and it is easy to accidentally make a change that prevents use of your site.

RELATED TOPICS

- Velocity Template Overview
- Customising Layouts
- Customising Look and Feel Overview
- Modify Confluence Interface Text

Customising the Login Page

It's fairly straightforward to customise the Confluence login page, to add your own logo or custom text. This will not customise the login *process* however, just what a user sees when she logs in.

> Customisations to the Confluence login page will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.

Only administrators with access to the server where Confluence is running can modify the Confluence login page.

Process to change the login page

1. Shut down your test instance of Confluence.

2. In the Confluence web application folder, find the file `confluence/login.vm`.

3. Make a copy of this file as a backup.

4. Edit the file with a text editor to make the required changes. The content is mostly HTML, but has some Velocity template variables in it. See Velocity Template Overview for more information about how these work.

5. Start Confluence up again and test your changes.

6. Apply the changes to your production Confluence instance.

The same process can be applied to modify most of the templates in the Confluence web application. Be careful to test your changes before applying them to a live site. The templates contain code that is vital for Confluence to function, and it is easy to accidentally make a change that prevents use of your site.

Related topics

- Editing the Global Logo
- Velocity Template Overview
- Customising Layouts
- Customising Look and Feel Overview
- Modify Confluence Interface Text

Styling Confluence with CSS

This page explains the facility for making visual changes to the look and feel of Confluence with CSS.

On this page:

- Introduction
- Considerations for Using Custom CSS
    - CSS Knowledge is Required
    - Security
    - Scaling
    - Features Cannot Be Disabled
    - Features Should Not Be Disabled
    - Confluence Version Compatibility

- - Test on Different Web Browsers
    - CSS Customisation is Not Supported
    - Follow the Tutorial
  - RELATED LINKS

Introduction

Cascading Style Sheets (CSS) are an industry-standard way of styling a web page. The content of a page is rendered with HTML, and its look and feel is determined by CSS files.

With the release of Confluence 2.10, you can easily upload a CSS text file and apply it to a space or even a whole Confluence instance. See this page for instructions.

> ✅ This function is not on by default. To allow space admins to edit stylesheets, go to the general configuration section of the admin console and turn on `"Enable Custom Stylesheets for Spaces"`.

Creating CSS styles that works seamlessly across different browsers is a delicate task for basic web sites, and reasonably challenging when customising web-applications like Confluence. It is important to carefully test each change that you make and ensure it works as expected in all areas of Confluence; for example, on the Confluence Dashboard as well as on regular pages.

In order to get you started, we have compiled this introduction, a basic styling tutorial and a more advanced tutorial which alters a Confluence menu.

Considerations for Using Custom CSS

CSS Knowledge is Required

If you're not familiar with CSS, this page has an accessible introduction. You should spend some time to become confident with Cascading Style Sheets before you start editing your Confluence instance.

Security

Custom CSS can be used to inject scripts into a page, opening the risk of cross-site scripting (XSS) attacks. With this feature enabled, space administrators could upload styles that steal other users' login credentials, trick their browsers into performing actions on the wiki without their knowledge, or even obtain global administration privileges. As such, this feature is disabled by default. Confluence administrators should only enable custom CSS if they are comfortable with the risks listed in this paragraph.

Scaling

Each page needs to scale. Depending on the resolution of the user's screen, the content should render intelligently. Your designs needs to degrade gracefully. Try resizing each page that exists in Confluence. There are quite a few pages in the browse-space-section, like drafts, labels, page hierarchy, and so on. Your style has to work everywhere, not just in the first page you happen to be looking at.

Features Cannot Be Disabled

It is easy to turn off certain links, headers, or even menu items by simply setting their style to 'hidden'. This can help you to roll out Confluence to users that may not be very Wiki-savvy yet. The simpler the UI, the easier it may be for them to use. However, please remember that removing the link to a part of the application does not mean that the functionality is not available. Every user can still change their style from within their browsers, or access the URL directly. Don't rely on CSS to disable parts of Confluence.

Features Should Not Be Disabled

Users familiar with Confluence will expect to find the same controls that they are accustomed to. Removing buttons or controls from the interface is not advised as it may frustrate your users and cause them to circumvent your design by using direct URL access, as mentioned above.

Confluence Version Compatibility

Be aware of any plans to upgrade your Confluence instance. Future versions of Confluence may not be compatible with your custom CSS — this may cause your CSS to break, requiring maintenance when Confluence is upgraded. Ask your Confluence administrator for more information.

Test on Different Web Browsers

As a rule you should test your modifications on different web browsers. Internet Explorer, Firefox, Opera and Safari (on Mac OS X) are some of the more popular browsers.

CSS Customisation is Not Supported

As creating custom CSS has potentially limitless possibilities, Atlassian will not support issues that are caused by or related to CSS customisation.

Follow the Tutorial

When you're ready, follow the examples in the Basic Styling Tutorial or the more advanced Styling Tabs in Confluence to get started.

RELATED LINKS

Basic Styling Tutorial
Styling Tabs in Confluence
Styling Fonts in Confluence
Including Cascading Stylesheets in Themes

Basic Styling Tutorial

This page contains instructions on how to get started with custom CSS styling in Confluence.

On this page:

- CSS Editing Quick-Start
- Tutorial: Changing the Header Background
- CSS Editing Tips
    - Begin With a Space Stylesheet
    - Use the Right Tools
    - Edit Simple Elements First
- RELATED LINKS

CSS Editing Quick-Start

1. Log in as the Space Administrator.
2. Open the Space Admin page.
3. Click '**Stylesheet**'.
4. Click '**Edit**' to change the code in the text field.
5. Paste your custom CSS into the text field. (Once text is saved here, you can make adjustments to it inline.)
6. Click '**Save**'. The new CSS will be visible on all content pages in the space.

Tutorial: Changing the Header Background

The header is the menu area at the top of a default Confluence page where the **Breadcrumb Links**, **Browse** menu, **User** menu and the **Quick Search** box reside. In this example, we are going to change the background of the header to include a custom graphic.

1. Create a custom graphic. For this example, we created a custom header graphic of 1046 x 61 pixels.
2. Upload the custom graphic to a page in the space that you are customising.
3. Note the page ID of the page where you uploaded the new graphic. (in this example, the page ID was '**658833839**'.
4. Compose your custom CSS for the header. The example below loads the new graphic (called '**header.png**') from a specific page (denoted by page ID '**658833839**') in the same space.

```
#header {
      background-image:url('../download/attachments/658833839/header.png');
      background-repeat: no-repeat;
    }
```

5. Log in as the Space Administrator.
6. Open the Space Admin page.
7. Click '**Stylesheet**'.
8. Click '**Edit**' to change the code in the text field.
9. Paste your custom CSS into the text field.
10. Click '**Save**'. Now reload the page (you may have to shift-reload). The background of the header will change.
11. The custom header will be visible on all content pages in the space. To revert your change, simple delete the custom code from the '**Stylesheet**' page and click '**Save**'.

CSS Editing Tips

Begin With a Space Stylesheet

A space stylesheet is a good starting point for CSS customisation, as it already includes all of the elements that can be changed. When you work on the space stylesheet it styles all content pages in the space. Build and test it at space-level, before considering applying the new stylesheet to your entire site. Once you are satisfied with your space design, test it thoroughly until you are confident that it has no problems. Then, you can look into advanced customisation of the Confluence CSS such as adjusting the Search page, the Dashboard and other integral pages.

Use the Right Tools

As the Confluence CSS is reasonably sophisticated, web development applications will help you to understand how the page styles have been created. In particular, you will need to view the existing source for the pages you're starting to work on. If you don't already have some, tools such as the following free applications will allow you to do this.

1. Firebug
Firebug, a plugin for the Firefox web browser, allows you to take a look at the style of each element on your page. This is very useful to see what styles are currently applied, for example styles applied to the header only.

2. Web Developer
The Web Developer plugin for Firefox allows you to edit CSS inline and create new page designs.

3. CSS Edit
CSS Edit is a stand-alone CSS editor for Macintosh that extracts all existing styles from a given page and allows you to overwrite these.

Edit Simple Elements First

Begin by editing simple elements and checking that they work. By making changes, then checking that each one worked, you can easily isolate any CSS code that is causing problems. Be aware that some page elements are more suited to customisation than others. For example, adding a gradient to the toolbar is less likely to 'break' the page than changing the page width. Editing reasonably static elements such as background graphics will render more predictably than designs which attempt to completely change the user interface or the Javascript-powered drop-down menus (which we don't recommend editing).

RELATED LINKS

Styling Confluence with CSS
Styling Tabs in Confluence
Including Cascading Stylesheets in Themes

Styling Fonts in Confluence

Confluence provides the ability to adjust its visual style via Cascading Style Sheets (CSS). With this feature, it is very easy to change the look and feel of Confluence.

This tutorial shows you to change the fonts of a normal Confluence page. We will show how to change the font and font sizes with a few lines of CSS.

Screenshot 1: Default Font in a Confluence Page
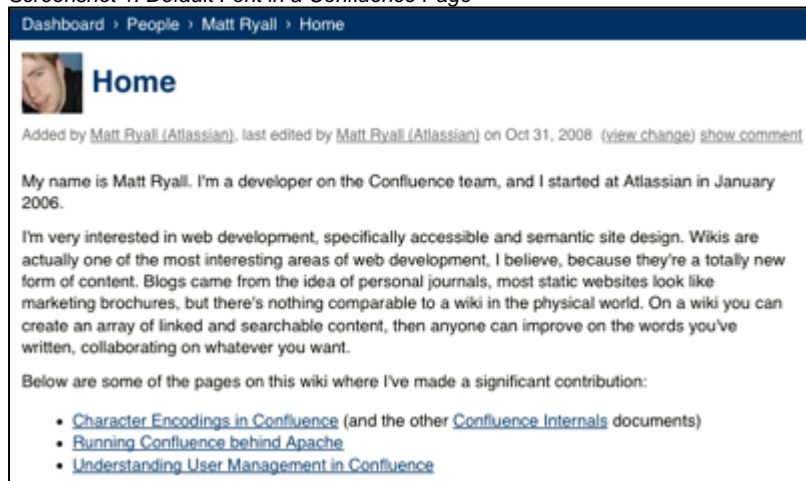
Screenshot 2: Custom Font in a Confluence Page

At the bottom of this page you can find the code for the custom font. Simply copy and paste it into the Space Stylesheet form within the Space Administration section.

Changing the fonts

In order to customise the fonts in Confluence, you first need to set the body font to the font you want. Second, you may want to adjust the font size to account for the fact that different fonts have different relative sizes.

The relevant CSS is shown below, and can be configured in Space Admin > Stylesheets. These styles change Confluence's font from its default of Helvetica/Arial – *sans serif* fonts – to Times/Times New Roman – *serif* fonts. To adjust for the fact that Times is a bit smaller than Helvetica, we increase the font size to 14 pixels.

```
01.body {
02.        font-family: Times, "Times New Roman", serif;
03.        font-size: 14px;
04.}
05..wiki-content,
06..wiki-content p,
07..wiki-content table,
08..wiki-content tr,
09..wiki-content td,
10..wiki-content th,
11..wiki-content ol,
12..wiki-content ul,
13..wiki-content li {
14.        font-size: 14px;
15.}
```

The many styles which include 'wiki-content' are necessary to change the font-size for all the tags in the wiki content.

RELATED LINKS

Basic Styling Tutorial
Styling Tabs in Confluence
Including Cascading Stylesheets in Themes

Styling Tabs in Confluence

Confluence provides the ability to adjust its visual style via Cascading Style Sheets (CSS). With this feature, it is very easy to change the look and feel of Confluence.

This tutorial shows you to change the look and feel of Confluence. We will address the Confluence tabs in this tutorial and how we can change their look completely with a few lines of CSS.

Let's take a look at what we are trying to achieve. Notice the dark blue tabs in *Screenshot 1*? Our goal is to blend them into the background and change the text style as in *Screenshot 2*.

Screenshot 1: Default Confluence Tabs in Admin Screen
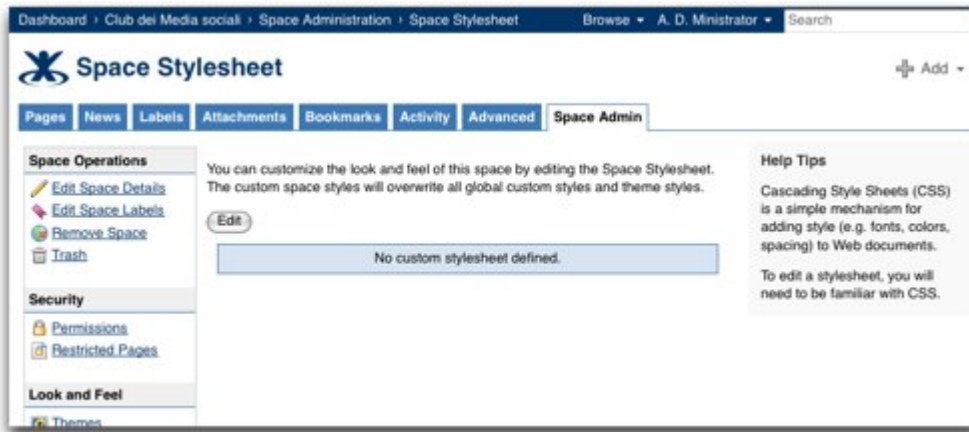
Screenshot 2: Customised Confluence Tabs in Admin Screen
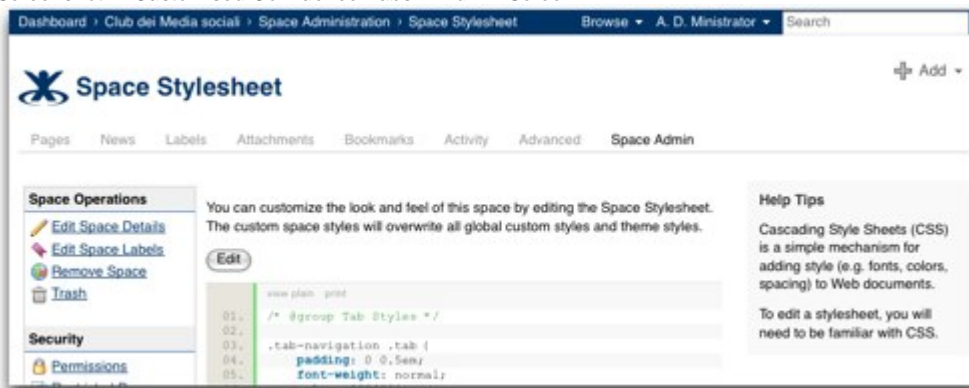
At the bottom of this page you can find the code for the custom tabs. Simply copy and paste it into the Space Stylesheet form within the Space Administration section. Let's discuss each selector in detail:

Container Style

```
1.#navigation, #tab-navigation{
2.        border-bottom: 1px solid #CCC;
3.        background-color: transparent;
4.        margin: 1em 0 2em -5px;
5.}
```

Lets start with the container around the tabs. We don't want a **background-color** for the tabs and the container, therefore we will set the color to be transparent. With the **border-bottom** attribute we will create a thin grey line at the bottom to separate the navigation from the rest of the page. Lastly we adjust the **margin**s of the container to create some space around the tabs.

Tab Style

```
1..tab-navigation .tab a {
2.      font-weight: normal;
3.      color: #999999;
4.      background-color: transparent;
5.      border: none;
6.}
```

Then we style each link within the tab. We set the **font-weight** to normal, to make the tabs less dominant and change the color to a dark grey. We also have to specify the **border** and **background-color** attributes explicitly to *overwrite* the default styles in Confluence. If we don't specify the background color for example, the blue color of the default style will be applied.

Hover Style

```
1..tab-navigation .tab a:hover {
2.      font-weight: normal;
3.      color: #0088CC;
4.      background-color: transparent;
5.      border: none;
6.}
```

Now we want to specify the **hover** attributes of the links. Note that we have to overwrite the **background-color** and the **border** attributes again, otherwise the default styles will be applied. We also change the color of the text for the hover effect of the link.

Currently Selected Element Style

```
1.ul.tab-navigation .current a:hover, ul.tab-navigation .current a {
2.    color:#000000;
3.    background-color: transparent;
4.    border:none;
5.}
```

The last selector we will need to customise is the element with the class current. Note that we use a more specific selector this time. The reason we are doing this is because these specific selectors are used in the default CSS stylesheet. If we were to use more general selectors, like we did above (ignoring the ul), the default style would still be applied since more specific selectors are rendered with a higher priority. We want to keep the general style of the other links but simply change the **color** of the text to black.

Complete CSS Style

Simply copy and paste the code below to change the look of the tabs in Confluence.

```
01./* @group Tab Styles */
02.
03..tab-navigation .tab a {
04.font-weight: normal;
05.color: #999999;
06.background-color: transparent;
07.border: none;
08.}
09.
10.
11..tab-navigation .tab a:hover {
12.font-weight: normal;
13.color: #0088CC;
14.background-color: transparent;
15.border: none;
16.}
17.
18.#navigation, #tab-navigation{
19.border-bottom-color:#CCC;
20.margin: 1em 0 2em -10px;
21.background-color: transparent;
22.}
23.
24.
25.ul.tab-navigation .current a:hover, ul.tab-navigation .current a {
26.               background-color: transparent;
27.               border:none;
28.               color:#000000;
29.}
30./* @end */
```

RELATED LINKS

Basic Styling Tutorial
Styling Tabs in Confluence
Including Cascading Stylesheets in Themes

Themes Configuration

- Applying a Theme To A Site
- Creating a Theme
    - Adding a theme icon
- Customising the Left Navigation Theme
- Deploying the theme as a plugin
- Including Cascading Stylesheets in Themes
- Modifying Look and Feel (for themes)
    - Configuring the theme plugin
- Themes Overview

Applying a Theme To A Site

Themes can can be applied across the site or to individual spaces.

Themes can be installed as plugins. Once a theme has been installed, a Confluence administrator can apply it via the **Administration**

Console. Once installed, themes become available to be applied across a site or to individual spaces. Any theme applied at the global level will become the default theme for all spaces in the site.

To apply a theme across the site,

1. Ensure that the theme plugin you wish to apply has been installed.

2. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**Themes**' under '**Look and Feel**' in the left navigation panel.
4. If there are any themes installed, they will be listed here.
5. Select a theme and click '**Confirm**'.

Screenshot : Applying a theme

Current Theme

The current theme controls the layout and colours of this space.

Default Theme
Assign the default Confluence look and feel. You can customise colour-schemes and layouts manually.

Choose New Theme

To change the theme of this space, select one below.

○ **Left Navigation Theme**
Provides a navigation bar on the left hand side of the screen.

○ **Clickr Theme**
A simple, clean Confluence theme inspired by Flickr - www.flickr.com.

Confirm

RELATED TOPICS

Adding a theme icon

Themes Overview

Including Cascading Stylesheets in Themes

Creating a Theme

Applying a Theme To A Site

Applying A Theme To A Space

Administrators Guide Home   Confluence Documentation Home

Creating a Theme

There are three steps involved involved in creating a theme:

- Modifying the look and feel of Confluence: Work with the different components that define the look and feel of Confluence and modify them to suit your theme:
  - Layout
  - Colour Scheme (optional)
  - Stylesheet (optional)

- Configuring the atlassian-plugin.xml file: Edit the central configuration file for the theme plugin to reference the new files defining your theme.
- Adding a theme icon: Add a preview icon for your theme.

- Deploying the theme as a plugin : Bundle the files into a jar file and deploy the theme as a plugin into Confluence.

Unsure what a theme is?

RELATED TOPICS

Adding a theme icon

Themes Overview

Including Cascading Stylesheets in Themes

Creating a Theme

Applying a Theme To A Site

Applying A Theme To A Space

Administrators Guide Home   Confluence Documentation Home

Adding a theme icon

A theme icon can be packed with a theme to give the user a little preview on how the theme will change the layout of Confluence. If you do not specify a custom icon for your theme, a default icon will be shown in the preview.

Defining the theme icon in the atlassian-plugin.xml

To include an icon in the theme, you will need to reference it as a [Downloadable Plugin Resource] from within the theme module.

Here is an example where an icon called my-theme-icon.gif is being used in the Dinosaur Theme:

```
01.<theme key="dinosaurs" name="Dinosaur Theme" class="com.atlassian.confluence.themes.BasicTheme">
02.        <description>A nice theme for the kids</description>
03.        <colour-scheme key="com.example.themes.dinosaur:earth-colours"/>
04.        <layout key="com.example.themes.dinosaur:main"/>
05.        <layout key="com.example.themes.dinosaur:mail-template"/>
06.        <resource name="themeicon.gif" type="download" location=
"com/example/themes/dinosaur/my-theme-icon.gif">
07.            <property key="content-type" value="image/gif"/>
08.        </resource>
09.</theme>
```

The resource parameter takes three arguments:

- **Name**: The name of the icon (⚠ has to be **themeicon.gif**).
- **Type**: The type of resource-in this instance, 'download'.
- **Location**: The location of the file represented in the jar archive you will use to bundle your theme.

The icon will automatically appear on the themes screen in the space and global administration and will be displayed next to the text and description of the theme.

○ **Default Theme**
Assign the default Confluence look and feel. You can customise colour-schemes and layouts manually.

Creating your own theme icon

In order to keep the look and feel of the icons consistent, we recommend to base the icon style on icons shipped with the Confluence themes. A good starting point when creating new icons is to use the default theme icon or the left navigation theme icon:

Unable to render embedded object: File (main-theme.gif! !leftnav-theme.gif) not found.

RELATED TOPICS

Adding a theme icon (Confluence Docs 3.1)

Themes Overview (Confluence Docs 3.1)

Including Cascading Stylesheets in Themes (Confluence Docs 3.1)

Creating a Theme (Confluence Docs 3.1)

Applying a Theme To A Site (Confluence Docs 3.1)

Applying A Theme To A Space (Confluence Docs 3.1)

Administrators Guide Home   Confluence Documentation Home

Customising the Left Navigation Theme

Introduction

Confluence comes bundled with the *Left Navigation* theme. This theme has a navigation menu on the left-hand side of the screen, which can be customised to contain additional links, sections and even macros.

Creating a custom navigation page

By default, the left-navigation theme just displays the space icon (or profile icon for a personal space), and three menus: page operations, browse space, and add content.

To add your own content to the top, create a page in your space called 'Navigation'. Put content there that you want to appear on the left navigation menu.

A couple of tips:

- items in a bulleted list show up as normal menu items
- use 'h1' to add a section heading for your menu items.

Examples

As an example, create a page called 'Navigation' with the following content:

```
h1. Search engines

* [Google|http://www.google.com]
* [Yahoo|http://www.yahoo.com]
* [MSN|http://search.msn.com]
```

This will give a left navigation menu like the image on the side of this page.

You can see another example of customised left-navigation theme on the Codegeist space with its associated Navigation page.
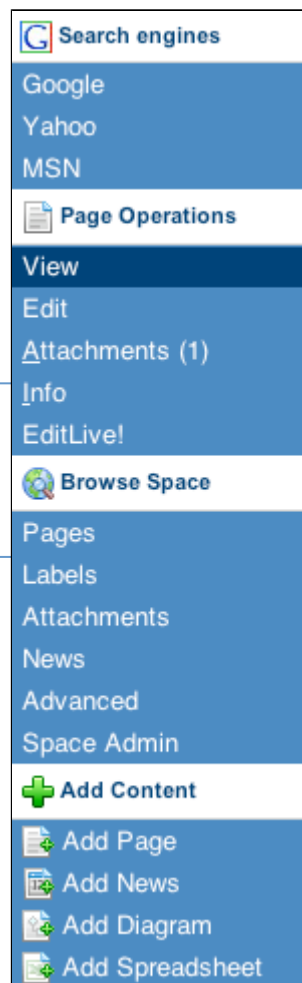
Custom left navigation

> ℹ️ **Alternative left-hand navigation**
> To insert an expandable/collapsible left-hand navigation menu, try the {pagetree} macro, supplied by the PageTree Plugin. You can follow the instructions to add the {pagetree} macro to your Confluence page layout. Note that this looks better if you use the Confluence Default theme rather than the Left Navigation theme.

RELATED TOPICS

Adding a Navigation Sidebar

Deploying the theme as a plugin

> ⚠️ This documentation is only applicable to **Confluence 2.5.x and earlier**. To create a new plugin in Confluence 2.6 and later, please refer to [Writing Confluence Plugins].

This page tells you how to deploy your own custom theme after you have created it. If you need to apply an already-installed Confluence theme to your Confluence site or space, please refer to one of the following pages instead:

- Applying a Theme To A Site
- Applying A Theme To A Space

In order to deploy your custom Confluence theme, you will have to have Ant installed. To learn how to install and use Ant, please follow the instructions on the projects website.

To deploy the theme, execute the following command from within the *theme* directory in your Confluence installation:

```
1.ant build -Dtheme=<specifynameoftheme>
```

For example to build a theme with the name *dinosaur,* you will have to type:

```
1.ant build -Dtheme=dinosaur
```

You will find the build jar of the dinosaur theme in you Confluence install directory under .../themes/dinosaur/dist_ directory.

Installing the theme

To install the theme you can simply drop the previously created .jar file into the `.../confluence/WEB-INF/lib` directory in your Confluence installation as described under Installing and Configuring Plugins.

As a second option, you can also call the following ant command instead of the one found above to install the theme and copy the jar automatically in the appropriate directory.

```
1.ant install -Dtheme=<specifynameoftheme>
```

Read more about plugins

To change the time and date formats,

> 1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
>
> 2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
>
>   - There are three time and date format settings:
>   - Time Format : determines the time format for when each news item is posted
>   - Date Time Format : determines date and time format for historical versions of pages.
>   - Date Format : determines date and time format for all new and modified content.
>
> 3. Change the formats using the guidelines in this document.
>
> 4. '**Save**' your changes.

RELATED TOPICS

📄 Adding a theme icon (Confluence Docs 3.1)

📄 Themes Overview (Confluence Docs 3.1)

📄 Including Cascading Stylesheets in Themes (Confluence Docs 3.1)

Administrators Guide Home   Confluence Documentation Home

Including Cascading Stylesheets in Themes

Confluence allows you to integrate your own stylesheets within the theme plugin so you can have greater control over the appearance of your site. Confluence's main stylesheet is a useful reference when overriding styles and can be found in the Confluence install directory under `...confluence/styles/site-css.vm`.

> **CSS for Confluence 2.6**
> Please refer to the information about changes in Confluence 2.6.

Step One: Defining the stylesheet in the atlassian-plugin.xml

To make a stylesheet available to a decorator, you will need to reference it as a resource from within the central configuration file - `atlassian-plugin.xml`.

Here is an example where a stylesheet is being used to define the 'leftnavigation' theme:

```
1.<layout key="main" name="Main Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
2.overrides="/decorators/main.vmd">
3.<resource type="velocity" name="decorator"
4.location="templates/leftnavigation/main.vmd"/>
5.<resource type="stylesheet" name="leftnav.css"
6.location="templates/leftnavigation/leftnav-css.vm">
7.</resource>
8.</layout>
```

The resource parameter takes three arguments:

- **Type**: The type of resource-in this instance, 'stylesheet'.
- **Name**: The name of the stylesheet.
- **Location**: The location of the file represented in the jar archive you will use to bundle your theme.

Step Two: Using the stylesheet in the decorator

To reference the stylesheet in the decorator, you will need to use the #pluginStylesheet velocity macro.

For example, here's how you reference the leftnav.css file defined in the layout entry above:

```
1.#pluginStylesheet("com.atlassian.confluence.themes.leftnavigation:main" "leftnav.css")
```

The macro takes two arguments:

- **completePluginKey**: The complete plugin key which is constructed from the pluginkey and the layout key like this: {**pluginKey**}:{**layoutKey**}
  In the above example, `com.atlassian.confluence.themes.leftnavigation` is the key of the plugin, and `main` is the key of the layout.
- **stylesheetName**: the name of the stylesheet

If you place your stylesheet **after** the `#standardHeader` macro in the decorator, the contents of your custom stylesheet will override those in Confluence's default stylesheet.

If your stylesheet needs to reference the colour scheme, you need to use the space stylesheet macro instead:

```
1.#pluginSpaceStylesheet("com.atlassian.confluence.themes.leftnavigation:main" "leftnav.css"
$spaceKey)
```

You can then use colour scheme references in your stylesheet, similar to Confluence's stylesheets, and they will be replaced with the appropriate global or space-specific colour scheme:

```
1..navItemOver {
2.        color: $action.navSelectedTextColor;
3.}
```

RELATED TOPICS

Adding a theme icon (Confluence Docs 3.1)

Themes Overview (Confluence Docs 3.1)

Including Cascading Stylesheets in Themes (Confluence Docs 3.1)

Creating a Theme (Confluence Docs 3.1)

Applying a Theme To A Site (Confluence Docs 3.1)

Applying A Theme To A Space (Confluence Docs 3.1)

Administrators Guide Home Confluence Documentation Home

Modifying Look and Feel (for themes)

Here's how you can define a new look and feel for Confluence in your theme:

1. **Layout** : Edit Confluence's layout by modifying the decorator files that are used to define it.
    - Working with Decorators
    - Velocity Template Overview
    - Configuring the atlassian.plugin.xml file to reference the decorators
2. **Colour schemes** : Configure a new colour scheme for your theme. **Optional**
    - Configuring a new colour scheme
    - Configuring the atlassian.plugin.xml file to include the new colour scheme
3. **Stylesheet** : Include a stylesheet to define your theme. **Optional**

 Note that for every component you edit, you will need to configure the `atlassian-plugin.xml` which is the central configuration file for the plugin to override the default files with the new files you've created.

Layout: Working with decorators

What are decorators?

Confluence is built on top of the Open Source SiteMesh library, a web-page layout system. To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a '.vmd' file and is written in a very simple programming language called **Velocity**. Learn more about Velocity.

Confluence comes bundled with a set of decorator or VMD files that you can customize. Broadly these are categorised into **Site**, **Content** and **Export** decorators. These are further grouped into categories called *contexts* and under each context has various *modes* (ways of viewing the context).

To make editing easier, layout for similar screens (example: view and edit page screens) is configured through the same VMD file. So, if you want to customize how the Confluence **View Page Screen** or **Edit Page Screen** looks, you can make **both** of these changes inside one decorator file: **page.vmd**.

Decorator	Context	Mode	Comment
page.vmd	page	'view', 'edit', 'edit-preview', 'view-information', and 'view-attachments'	
blogpost.vmd	blogpost (news)	'view', 'edit', 'edit-preview', and 'remove'	We prefer to use 'news' as an end-user term; all templates and classes use 'blogpost' to indicate RSS related content
mail.vmd	mail	'view', 'view-thread' and 'remove'	

space.vmd	space-pages, space-mails, space-blogposts, space-templates, space-operations, space-administration	CONTEXT: "space-pages". MODES: "list-alphabetically", "list-recently-updated", "list-content-tree", "create-page". CONTEXT: "space-mail". MODES: "view-mail-archive". CONTEXT: "space-blogposts". MODES: "view-blogposts", "create-blogpost". CONTEXT: "space-templates". MODES: "view-templates". CONTEXT: "space-operations". MODES: "view-space-operations". CONTEXT: "space-administration". MODES: "view-space-administration", "list-permission-pages".	space.vmd handles a wide range of options, this context is accessed by clicking on 'browse space' in the default theme of Confluence (tabbed theme)
global.vmd	global	'dashboard', 'view-profile', 'edit-profile', 'change-password-profile', 'edit-notifications-profile'	
main.vmd	n/a (header and footer formatting)		main.vmd is used to control the header and footer of each page, not the page specific presentation logic

For example, if you wanted to remove the 'Attachments' tab on the view page screen, you would make this layout change in the page.vmd file - where the 'view' mode is handled (as shown below).

```
01.#*
02.    Display page based on mode: currently 'view', 'edit', 'preview-edit', 'info' and
'attachments.
03.    See the individual page templates (viewpage.vm, editpage.vm, etc.) for the setting of the
mode parameter.
04.*#
05.## VIEW
06.#if ($mode == "view")
07.
08.          <make layout modifications here>
09.
10.#elseif ...
```

Step One: Copying the decorators

The easiest way to begin configuring a new layout is by copying the default decorator files and editing them to suit your theme.

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select **Layouts** in the left panel. This will display options to view and edit the default decorators.
3. Copy the files that you intend to modify and place them in a directory structure that makes sense to you. See example below.

Step Two: Creating a directory structure for the decorators:

You should place your decorators in a directory hierarchy which makes sense to you. We recommend that you place the atlassian-plugin.xml file at the top level of the directory structure, and then place the decorators in directories which make a meaningful division of what they do.

Here is an example:

```
1.atlassian-plugin.xml
2.com/atlassian/confluence/themes/mytheme/
3.com/atlassian/confluence/themes/mytheme/global.vmd
4.com/atlassian/confluence/themes/mytheme/space.vmd
5.com/atlassian/confluence/themes/mytheme/mail.vmd
6.com/atlassian/confluence/themes/mytheme/blogpost.vmd
7.com/atlassian/confluence/themes/mytheme/main.vmd
8.com/atlassian/confluence/themes/mytheme/page.vmd
```

Step Three: Editing the decorators

To edit the decorators, you will require knowledge of a very simple programming language called **Velocity**. Learn more about Velocity.

Decorator Macros

When editing the decorators, you will need to use **Decorator Macros** to draw complex or variable parts of the page such as menus and breadcrumbs. See Working With Decorator Macros

Theme Helper Object

When editing decorator files you will also come across a variable called **$helper** - this is the **theme helper object**.

The following table summarises what this object can do:

Behaviour	Explanation
$helper.domainName	displays the base URL of your Confluence instance on your page. This is useful for constructing links to your own Confluence pages.
$helper.spaceKey	returns the current space key or null if in a global context.
$helper.spaceName	returns the name of the current space
$helper.renderConfluenceMacro("{create-space-button}")	renders a call to a Confluence Macro for the velocity context
$helper.getText("key.key1")	looks up a key in a properties file matching ` 1.key.key1=A piece of text` and returns the matching value ("A piece of text")
$helper.action	returns the XWork action which processed the request for the current page.

If you are on a page or space screen you also have access to the actual page and space object by using `$helper.page` and `$helper.space` respectively.

If you want to deliver more into what other methods are available in this object, please see our API's for ThemeHelper.

Step Four: Configuring the central configuration file to reference the new decorators

How to do this is explained in Configuring the theme plugin

Working with colour schemes for themes

Configuring the colour scheme

The easiest way to configure a colour scheme is to do it dynamically from the **Administration Console** (as you would normally when you want to change the site's colour scheme online), and then express it as an xml file. This method makes it possible for you to experiment with different colours and test them out before including the colour scheme in your theme.

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Colour scheme**' in the left panel.
3. Use the colour picker to define the colours for the following UI elements:

- **Top Bar** - the bar across the top of the page that contains the breadcrumbs.
- **Space Name Text** - the text of the current space name located above the page title.
- **Heading Text** - all heading tags throughout the space.
- **Links** - all links throughout the space.
- **Borders and Dividers** - table borders and dividing lines.
- **Menu Bar Background** - background of top navigational buttons
- **Menu Bar Text** - text that appears on the menu bar
- **Menu Bar Background Highlight** - background colour of menu bar when highlighted.
- **Menu Bar Text Highlight** - menu bar text when highlighted

More information on customising colour schemes

Expressing the colour scheme as XML

Once, you have decided on the colours for the different UI elements, you will need to configure the atlassian.plugin.xml to include the new colour scheme. How to do this is explained in detail in Configuring the theme plugin.

RELATED TOPICS

Adding a theme icon

Administrators Guide Home Confluence Documentation Home

Configuring the theme plugin

Each plugin is described in its own atlassian-plugin.xml file, which specifies attributes of the plugin, including a description of each module it contains. Once you have modified the different components to define a new look and feel for your theme, you will need to configure this file so Confluence knows where to look when overriding the default files.

The easiest way to begin is by copying the atlassian-plugin.xml from one of the default themes bundled with confluence and modifying it for your theme.

The structure of an atlassian-plugin.xml file is fairly self-explanatory. In the code segment below you will find a full example of an atlassian-plugin.xml:

```
01. <atlassian-plugin key="com.atlassian.confluence.themes.tabless" name="Plain Theme">
02.         <plugin-info>
03.                 <description>This theme demonstrates a plain look and feel for Confluence. It
is useful as a building block for your own themes.</description>
04.                 <version>1.0</version>
05.                 <vendor name="Atlassian Software Systems Pty Ltd" url="
http://www.atlassian.com/"/>
06.         </plugin-info>
07.
08.         <theme key="tabless" name="Tabless Theme" class=
"com.atlassian.confluence.themes.BasicTheme">
09.                 <description>plain Confluence theme.</description>
10.                 <layout key="com.atlassian.confluence.themes.tabless:main"/>
11.                 <layout key="com.atlassian.confluence.themes.tabless:global"/>
12.                 <layout key="com.atlassian.confluence.themes.tabless:space"/>
13.                 <layout key="com.atlassian.confluence.themes.tabless:page"/>
14.                 <layout key="com.atlassian.confluence.themes.tabless:blogpost"/>
15.                 <layout key="com.atlassian.confluence.themes.tabless:mail"/>
16.                 <colour-scheme key="com.atlassian.confluence.themes.tabless:earth-colours"/>
17.         </theme>
18.
19.         <layout key="main" name="Main Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
20.                         overrides="/decorators/main.vmd">
21.                 <resource type="velocity" name="decorator"
22.                                         location=
"com/atlassian/confluence/themes/tabless/main.vmd"/>
23.         </layout>
24.
25.         <layout key="global" name="Global Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
26.                         overrides="/decorators/global.vmd">
27.                 <resource type="velocity" name="decorator"
28.                                         location=
"com/atlassian/confluence/themes/tabless/global.vmd"/>
29.         </layout>
30.
31.         <layout key="space" name="Space Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
32.                         overrides="/decorators/space.vmd">
33.                 <resource type="velocity" name="decorator"
34.                                         location=
"com/atlassian/confluence/themes/tabless/space.vmd"/>
35.         </layout>
36.
37.         <layout key="page" name="Page Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
38.                         overrides="/decorators/page.vmd">
39.                 <resource type="velocity" name="decorator"
40.                                         location=
"com/atlassian/confluence/themes/tabless/page.vmd"/>
41.         </layout>
42.
```

```
43.         <layout key="blogpost" name="Blogpost Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
44.                            overrides="/decorators/blogpost.vmd">
45.              <resource type="velocity" name="decorator"
46.                                 location=
"com/atlassian/confluence/themes/tabless/blogpost.vmd"/>
47.         </layout>
48.
49.         <layout key="mail" name="Mail Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
50.                            overrides="/decorators/mail.vmd">
51.              <resource type="velocity" name="decorator"
52.                                 location=
"com/atlassian/confluence/themes/tabless/mail.vmd"/>
53.         </layout>
54.
55.
56.         <colour-scheme key="earth-colours" name="Brown and Red Earth Colours"
57.                                       class=
"com.atlassian.confluence.themes.BaseColourScheme">
58.               <colour key="topbar" value="#440000"/>
59.               <colour key="spacename" value="#999999"/>
60.               <colour key="headingtext" value="#663300"/>
61.               <colour key="link" value="#663300"/>
62.               <colour key="border" value="#440000"/>
63.               <colour key="navbg" value="#663300"/>
64.               <colour key="navtext" value="#ffffff"/>
65.               <colour key="navselectedbg" value="#440000"/>
66.               <colour key="navselectedtext" value="#ffffff"/>
67.         </colour-scheme>
68.
69.</atlassian-plugin>
```

Modifying the {{atlassian-plugin.xml}}file

We will configure this file section by section.

Plugin information

```
1.<atlassian-plugin key="com.atlassian.confluence.themes.tabless" name="Plain Theme">
2.     <plugin-info>
3.             <description>This theme demonstrates a plain look and feel for Confluence. It
is useful as a building block for your own themes.</description>
4.             <version>1.0</version>
5.             <vendor name="Atlassian Software Systems Pty Ltd" url="http://www.atlassian.com/
"/>
6.     </plugin-info>
```

plugin key : Specify a key that uniquely identifies the plugin, eg. com.example.themes.dinosaur

name : Give the plugin a name.

description : Provide a short description of the plugin.

vendor : Replace the text with your information.

Theme information

```
1.<theme key="dinosaurs" name="Dinosaur Theme"
2.              class="com.atlassian.confluence.themes.BasicTheme">
3.        <description>A nice theme for the kids</description>
4.        <colour-scheme key="com.example.themes.dinosaur:earth-colours"/>
5.        <layout key="com.example.themes.dinosaur:main"/>
6.        <layout key="com.example.themes.dinosaur:mail-template"/>
7.</theme>
```

theme key : Specify a key that uniquely identifies the theme.

class : The class of a theme must implement com.atlassian.confluence.themes.Theme. The com.atlassian.confluence.themes.BasicTheme class provided with Confluence gathers together all the resources listed within the module definition into a theme.

name : Give the theme a name. Make sure that you replace all instances of the theme name with this name.

description : Provide a short description of your theme

colour-scheme key : A theme can contain an optional colour-scheme element that defines which colour-scheme module this theme will use. If you are using a new colour scheme, enter its key.

layout key : A theme can contain any number of `layout` elements that define which layouts should be applied in this theme. Refer to these modules by their **module complete key** as shown above.

Referencing the decorators

You will need to add a layout entity as shown below for each of the decorators you are using. See working with decorators

```
1.<layout key="page" name="Page Decorator" class="com.atlassian.confluence.themes.VelocityDecorator"
2.                              overrides="/decorators/page.vmd">
3.              <resource type="velocity" name="decorator"
4.                                  location=
"com/atlassian/confluence/themes/tabless/page.vmd"/>
5.</layout>
```

class : The class which each decorator, or layout, is mapped to must implement `com.atlassian.confluence.themes.VelocityDecorator`.

overrides : The layout entry must provide an `overrides` attribute which defines which decorator within Confluence is being overrridden by the theme.

Location : Specify the location of the new decorator file, so Confluence know where to look when overriding the default decorator.

🛈 It is possible for a theme to use modules that aren't in the same plugin as the theme. Just keep in mind that your theme will be messed up if the plugin that the theme depends on is removed.

Including the colour scheme

Colour schemes can be pre-configured for your theme dynamically from the **Administration Console**. See configuring colour schemes

To transport them within a theme however, they need to be expressed in the `atlassian-plugin.xml` file as shown above.

```
01.<colour-scheme key="earth-colours" name="Brown and Red Earth Colours"
02.                               class="com.atlassian.confluence.themes.BaseColourScheme"
>
03.          <colour key="topbar" value="#440000"/>
04.          <colour key="spacename" value="#999999"/>
05.          <colour key="headingtext" value="#663300"/>
06.          <colour key="link" value="#663300"/>
07.          <colour key="border" value="#440000"/>
08.          <colour key="navbg" value="#663300"/>
09.          <colour key="navtext" value="#ffffff"/>
10.          <colour key="navselectedbg" value="#440000"/>
11.          <colour key="navselectedtext" value="#ffffff"/>
12.</colour-scheme>
```

colour-scheme key : Specify a key that uniquely identifies the colour scheme.

name : Give a name to the colour scheme.

class : The class of the colour scheme must implement `com.atlassian.confluence.themes.ColourScheme`. The `com.atlassian.confluence.themes.BaseColourScheme` class provided with Confluence sets the colours based on the module's configuration.

colour key: For each UI element, you will need to add its name and value.

see configuring colour scheme

RELATED TOPICS

📄 Adding a theme icon

📄 Themes Overview

📄 Including Cascading Stylesheets in Themes

📄 Creating a Theme

📄 Applying a Theme To A Site

📄 Applying A Theme To A Space

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Themes Overview

Themes are pre-defined styles that can be applied to alter the appearance of your site.

Use themes when you want to add new functionalities or to change the appearance of Confluence. For example, you will need to use themes to apply a left-navigation scheme instead of the default top-navigation scheme.

Themes are installed as plugins and added via the **Administration Console**. Once installed, themes can be applied across the site or to individual spaces.

To look at the themes installed,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select 'Themes' under 'Look and Feel' in the left navigation panel.
3. If there are any themes installed, they will be listed here.

What do you want to do?

Apply a theme

Create a new theme

Include cascading stylesheets in a theme

RELATED TOPICS

Adding a theme icon

Themes Overview

Including Cascading Stylesheets in Themes

Creating a Theme

Applying a Theme To A Site

Applying A Theme To A Space

Administrators Guide Home   Confluence Documentation Home

Importing Data

- Importing Content from another Wiki
- Snip Snap Import

- Universal Wiki Converter
- Importing Content Into Confluence

Importing Content from another Wiki

The Universal Wiki Converter (UWC) allows you to import content from other wikis into Confluence. The Confluence Administration Console offers a link to the Universal Wiki Converter documentation and download sites.

> **You need to install and run the UWC separately from Confluence.**

The UWC is a standalone application that communicates with Confluence remotely. You cannot install the UWC directly into Confluence. Instead, download the UWC separately and run it according to the instructions below.

The UWC supports many wiki dialects. In addition, the UWC is an extensible framework, which means that developers can continue writing new conversion modules for other wikis. To see the latest list of conversions available, please refer to the UWC documentation.

- Download the latest version of the UWC.
- For information on installation and usage, see the UWC Quick Start Guide.
- For information on developing your own converter module, see the UWC Developer Documentation.

- For information about a specific wiki, including a list of currently supported wikis, see the UWC documentation.
- If you have questions or would like to share information about the UWC, please visit the UWC Forum.

Screenshot: Links from the Confluence Administration Console to the UWC

RELATED TOPICS

- Importing Content Into Confluence
- Importing Data

Snip Snap Import

The snipsnap importer allows you to import a Snip Snap XML backup file into a space in Confluence.

> **What is Snip Snap?**
> Snip Snap is a wiki used as a knowledge and content management tool. For more information, take a look at the Snip Snap home page and the Wikipedia page about Snip Snap.

Some limitations:

- Currently, attachments are not imported, and Confluence does not recognise duplicate users.
- You cannot import content into multiple spaces.

You need to have System Administrator permissions in order to perform this function.

To import a Snip Snap backup file into Confluence,

1. First, use Snip Snap to export a backup to an XML file. Now return to Confluence.

2. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**SnipSnap Import**' in the left-hand panel.
4. Enter the location of the Snip Snap backup file in the input field displayed. You can also '**browse**' and locate the file.
5. Select a space to import the content into and click '**Save**'.

RELATED TOPICS

Importing Data

Administrators Guide Home   Confluence Documentation Home

Installing Plugins and Macros

A **plugin** is an add-on to the core Confluence code, which can extend the Confluence functionality. Some plugins are shipped with Confluence, others are available for you to install yourself.

A **macro** allows a developer to perform programmatic functions within a page, and gives the Confluence user access to more complex content structures. Many macros are made available by plugins.

Installing Plugins

You need to have System Administrator permissions in order to install and configure plugins. This page introduces two methods of installing plugins:

- Via the Plugin Repository Client
- Manually

Installing and Configuring Plugins using the Plugin Repository Client

If the plugin you wish to install is listed in the Confluence Plugin Repository, you can use the Confluence Repository Client to install it. In the 'Administration' section of Confluence, click the Plugin Repository. Then find the plugin in the list and click the 'Install' link. There's more information in Installing and Configuring Plugins using the Plugin Repository Client.

Installing and Configuring Plugins Manually

If the plugin you wish to install is not listed in the Confluence Plugin Repository, you can still install it by uploading the plugin jar file to your Confluence site. In the 'Administration' section of Confluence, click the 'Plugins' link.

There's more information in Installing and Configuring Plugins Manually.

Troubleshooting

Problem Loading the Plugin Repository with IE7

There is a problem using the Plugin Repository with Internet Explorer 7. You may see an error message saying 'Object error'. In this case, the browser will just hang until you close the browser window or tab.

This problem is caused by a bug in IE7, which Microsoft have fixed in Windows Service Pack 3. Details are on our JIRA site at CONF-10837.

There are some workarounds for those who cannot install the Microsoft Service Pack. You can use a different browser instead of IE7, such as Firefox. Or you can upload the plugin jar into Confluence manually, via the 'Plugins' option in the Administration Console. See Installing and Configuring Plugins Manually.

RELATED TOPICS

- Installing and Configuring Plugins Manually
- Installing and Configuring Plugins using the Plugin Repository Client
- Plugin loading strategies in Confluence
- Removing Malfunctioning Plugins
- Enabling and Configuring Macros
  - Configuring a URL Whitelist
  - Configuring the userlister Macro

Installing and Configuring Plugins Manually

On this page:

- Installing a Plugin Manually
- Enabling and Disabling Plugins

This document is for **administrators** who wish to install new plugins or manage the plugins installed in their Confluence server. For an overview of how plugins work in Confluence, read the Confluence Plugin Guide.

You need to have System Administrator permissions in order to install and configure plugins.

Looking for existing plugins?
See the existing plugins and extensions written by the community in the Confluence Extensions space.

The Plugin Repository
The Plugin Repository provides an alternative way to install plugins directly from the the plugin libraries.

Plugin Safety
Plugins are very powerful: they can change the behaviour of almost any part of the Confluence server. This makes it **very important** that you trust a plugin before you install it. Always be aware of where (and who) a plugin comes from.

Installing a Plugin Manually

Plugins are distributed as a jar file. To install a plugin:

1. In the 'Administration' section of Confluence, click the 'Plugins' link.
2. Use the 'Choose file' button to find the plugin jar you wish to install from your hard drive or network location, and select it.
3. Click 'Upload'.
4. The plugin will be uploaded to Confluence and will be automatically installed.
5. Check the 'Plugin Administration' screen to ensure if the plugin is available.
6. Enable the plugin if necessary. (Some plugins will be enabled by default when they are installed. Others will have to be manually enabled from the Plugin Administration screen.)

Enabling and Disabling Plugins

As administrator, you can enable and disable plugins, and the plugin modules which form part of each plugin. You can do this from the 'Plugins' section of the administration screen. All plugins installed in the Confluence server are listed on the left-hand side. To enable or disable a plugin or its modules, click the plugin name.

On the right-hand side, a description of the plugin is shown, including its component plugin modules.

You can enable or disable the whole plugin:

Or each module individually:

⚠ Disabling a plugin module *may* cause other modules in the same plugin to cease to function correctly. When in doubt, make sure you disable or enable the entire plugin.

RELATED TOPICS

Installing and Configuring Plugins using the Plugin Repository Client
Installing Plugins and Macros

Installing and Configuring Plugins using the Plugin Repository Client

⊖ **Plugin Safety**
Plugins are very powerful: they can change the behaviour of almost any part of the Confluence server. This makes it very important that you trust a plugin before you install it. To help ensure the authenticity of a plugin, always be aware of its origins and/or its creator.

The Plugin Repository provides an easy way to install and configure plugins.

- If you are using Confluence version 2.3 or later, the Plugin Repository is included as part of Confluence.
- If you are using Confluence 2.0 - 2.2.10, you can use the Plugin Repository after you install the Confluence Repository Client.
- Confluence versions prior to 2.0 cannot use the Plugin Repository, so you will need to install and configure plugins manually instead.

ℹ You need to have System Administrator permissions in order to install and configure plugins.

On this page:

- Using the Plugin Repository
    - Filtering the List of Plugins
    - Using the List of Plugins
- Viewing Plugin Modules and Versions
    - Administering the Plugin Repository
    - Uploading a Non-Repository Plugin
    - Obtaining More Information About Plugins in the Plugin Repository
- Troubleshooting

Using the Plugin Repository

Go to the '**Administration Console**' and click '**Plugin Repository**' in the left-hand panel. The following will be displayed:

Screenshot: Plugin Repository

Filtering the List of Plugins

Along the top of the page, you'll see three items which allow you to choose the plugins you want displayed:

- **Status filter** — Defaults to 'All Plugins'. Choose one of the following to limit the list of plugins displayed:
    - 'Installed Plugins' – Plugins which have been installed on your Confluence instance.
    - 'Outdated Plugins' — Plugins for which updates are available
    - 'Available Plugins' — Plugins which are available, but have not yet been installed.
    - 'Supported Plugins' — All plugins which are supported by Atlassian or a third-party.
- **Search** — Enter text in the middle textbox to quickly find what you're looking for:
    - Quickly locate plugins by searching on their title, description, vendor and other details. Just type what you are looking for and let Confluence do the rest

- The search results are filtered by the status filter (as discussed above), so if you want to see all installed plugins from Adaptavist.com, for example, set the filter to 'Installed Plugins' and type 'Adaptavist.com' into the search box.
- **Categories list** — Filter by category:
  - Simply choose the desired category from the list and only plugins relating to that category will be displayed.
  - Note that plugins can exist in more than one category.

Using the List of Plugins

Under the filter options, the list of plugins matching the current filter settings is shown in a table.

Each column in the table shows information about a particular plugin, and allows you further configuration options:

- **Plugin Name** — Displays the name of the plugin (linked to the detailed information page)and the plugin vendor (linked to their website if applicable)
  - Click the '**+**' icon to expand the information display showing plugin description.
  - Click the '**-**' icon to hide the description again.
- **Payment** — Can be one of the following:
  - Free (self explanatory)
  - Donate (it's free, but you should consider donating to keep it that way)
  - Buy (it's commercial - click the link to show a price list and purchase online)
- **Status** — Shows the current status of this plugin in respect to your Confluence installation:
  - Installed - installed and up-to-date
  - Outdated - installed, but there are new versions available
  - Available - not installed yet
  - Non Repository - a version is installed which is not in the repository
- **Support** — Tells you who supports the plugin:
  - 'Atlassian' — The plugin is supported by Atlassian. If you have any problems, please raise a ticket at the Atlassian Support System.
  - 'Unsupported' — The plugin has been developed by a third party, not by Atlassian, and is not currently supported by Atlassian. In addition, the third party has not yet given detailed information about support arrangements. This does **not** necessarily mean that the plugin is not supported. Please refer to the plugin's home page in the Confluence Extension space or the Atlassian Plugins Exchange site.
  - The 'Support' column can also contain a link to the third-party plugin support site.
  - Read more information about supported plugins.
- **Install** — Install, upgrade or uninstall a plugin:
  - When installing or upgrading, everything is automatic (i.e. it downloads and installs for you, etc). Although the client (since 1.0.2) warns you of dependencies and (since 1.0.3) will do its best to check what has been downloaded is what you asked for - Confluence may break as if you had uploaded the plugins to the Plugin Manger yourself. Where it can, the client will error constructively allowing you to choose the best course of action for yourself. In general, **things usually work - and if they don't its a bug with the client or the metadata**.
  - If the Confluence Repository Client encounters a password request when downloading the plugin (usually case with Commercial plugins), you will be prompted for a username and password.
  - If the plugin is installed into WEB-INF (or otherwise uninstallable) it will display 'Manually Installed. where the actions would be.
- **Enabled** — If the box is ticked, the plugin is enabled, otherwise it's either disabled or partially disabled. You can enable or disable individual modules within the plugin from the plugin details screen (see later).
- **Configure** — If the plugin offers further configuration options, you can click the 'Configure' link. A new screen will open, showing the specific options offered by the plugin.

You can click the table headings to sort the table. Click a second time to reverse the sort.

Viewing Plugin Modules and Versions

When you click the name of a plugin in the plugin list, you'll be taken to the detailed view for that plugin.

General information and plugin module details are shown at the top of the display and from here you can disable or enable individual modules.

Note: Disabling a plugin module *may* cause other modules in the same plugin to cease to function correctly. When in doubt, make sure you disable or enable the entire plugin.

Near the bottom of the display a table outlines all plugin versions and shows which you have installed. Just like the plugins list, you can click the '**+**' to expand the details shown for a specific version.
This screen also allows you to quickly install, upgrade, downgrade and uninstall any version of the plugin.

Administering the Plugin Repository

There are various settings on the '**Admin**' tab.

The most important of these is the **'Data Source'** — without this, you'll see no plugins in the list and will get a fair number of errors.

The **'Earliest Plugin State'** allows you to filter the plugin list to versions at or above a specific state: Alpha, Beta, Release Candidates, Stable. If you are running in a production environment, you will usually want to set this setting to 'Stable + Release Candidates' or 'Stable Only'.

The **'Plugin Compatibility'** setting allows you to restrict the list to only show plugin versions that are specifically known to work with your version of Confluence.

ℹ In many cases, plugins will work with your version of Confluence, but they might be marked as 'unknown' compatibility if the plugin author hasn't been able to test with that specific Confluence version. As such, it's extremely useful if you find an 'unknown' version of a plugin to work (or not work) with your Confluence that you let us know (by any means possible) so that we can update the repository to reflect this.

The **'Category Visibility'** setting allows you to trim down the categories list by hiding categories that don't contain a plugin yet.

Uploading a Non-Repository Plugin

There are several plugins that are not currently listed by the Plugin Repository which need to be manually uploaded either as a file or from a remote server URL.

If you want to install such a plugin:

1. Click the '**Upload**' tab.
2. Enter either the file name or the URL of the remote server URL.
3. Click the 'Install' button.

Obtaining More Information About Plugins in the Plugin Repository

Above we have described the **'Plugin Repository'** screen in the **Administration Console**. For more information about these plugins, please refer to its page on the Atlassian Plugins Exchange site.

ℹ To find a specific plugin quickly on this site, ensure that you are in the **Plugins** section of this site and use the search tool. You should only need to enter a few keywords of the plugin's name to find it.

Troubleshooting

Problem Loading the Plugin Repository with IE7

There is a problem using the Plugin Repository with Internet Explorer 7. You may see an error message saying 'Object error'. In this case, the browser will just hang until you close the browser window or tab.

This problem is caused by a bug in IE7, which Microsoft have fixed in Windows Service Pack 3. Details are on our JIRA site at CONF-10837.

There are some workarounds for those who cannot install the Microsoft Service Pack. You can use a different browser instead of IE7, such as Firefox. Or you can upload the plugin jar into Confluence manually, via the 'Plugins' option in the Administration Console. See Installing and Configuring Plugins Manually.

RELATED TOPICS

Confluence Plugin Guide

Plugin loading strategies in Confluence

The categories

Confluence plugins have different behaviour based on how they are loaded by Confluence. The plugins themselves are the same, but based on how they are loaded, they may or may not be upgraded, or may not be disabled, or may not be uninstalled. This chart should explain how plugins can be loaded by Confluence, and the ramifications for each choice.

The category any *particular* plugin is in can vary with Confluence version or circumstance. The examples mentioned here describe the way particular plugins are loaded by default in Confluence 2.8.

Category	Description	Example
Static	**cannot be installed or upgraded without a Confluence restart**	
Core	Included with Confluence and cannot be uninstalled. The classes and plugin.xml are not bundled into plugin jars, but mixed in with Confluence source on the main classpath. Additionally, the plugin.xml definitions are not called "atlassian-plugin.xml" as they are everywhere else, but are named for the plugin e.g., "basic-macros.xml". We would like to separate some of them out and turn them into *Bundled* plugins.	Admin Sections
WEB-INF/lib	Confluence also places some plugin jars inside **WEB-INF/lib**. They are inserted during the build process by Maven. These plugins, likewise, cannot be uninstalled. In ancient times, this was the only way to install plugins, so users are also free to install plugins here. We try to discourage them from doing so, however. As of version 3.0, most of the JAR files in this directory are library dependencies, not plugins.	
Dynamic	**the opposite of static, these can be installed/upgraded while Confluence is running**	

Bundled	Bundled plugins can be administered from the Plugins console from Administration >> Plugins. You can upload or disable them there.	Office Connector
	Bundled plugins are included in a zip of jars called **atlassian-bundled-plugins.zip** which is on the main Confluence classpath, in a resources directory - `<confluence-install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup`. At Confluence startup, they are extracted and copied into the **$CONFLUENCE_HOME/bundled-plugins** directory, from whence they are loaded. To remove a bundled plugin (you shouldn't normally have to do this), remove the plugin from the atlassian-bundled-plugins.zip file and the bundled-plugins directory, otherwise Confluence will just put it back in place on the next startup. In versions later than 2.6, you'll have to recreate the .jar file (if the jar file is from the lib folder) or recreate the zip folder(if its in the classes folder). Bundled plugins can be upgraded or disabled.	
Uploaded	Installed by the user via the plugin repository or the Plugin Manager page. These plugins are stored in the database and then copied to the **$CONFLUENCE_HOME/plugins-cache** folder on each Confluence node.	could be anything

To summarise the relationships of categories in the table, all plugins are either *Static* or *Dynamic*. *Static* plugins can be further categorised into *Core* or *WEB-INF/lib*. *Dynamic* plugins are divided into *Bundled* and *Uploaded*.

Use of the categories in Confluence

Within Confluence, the *Core* and *WEB-INF/lib* categories are not actually named as such, and they don't map neatly to other names (though they do map, as will be explained). They are used here because of the logical distinction they provide.

In Confluence, some of the *Core* plugins are called "System". Plugins can be designated as "System" by adding a flag to the plugin manifest file. To do this, `system=true` should be added to the top-level `atlassian-plugin` element of the manifest file. The manifest file is generally called `atlassian-plugin.xml`, but it could have another name; the *Core* plugins' files do.

All of the *Core* plugins once were labeled as "System", but it seems the practice has faded over time. If a plugin is designated as "System", then it will not show up in the Plugin Manager page in Confluence and thus cannot be enabled/disabled. However, it *will* show up in the Plugin Repository Client, where it can be disabled; allowing disabling there is probably incorrect behavior.

Static plugins that are not marked as "System" (any remaining *Core* and *WEB-INF/lib* plugins), are simply called *Static* in Confluence. There is no way to tell the *WEB-INF/lib* and *Core* plugins apart from within Confluence. You just have to figure out where the classes are.

Members of the other specific categories - *Bundled* and *Uploaded* - can be determined. We can tell which plugins are *Bundled* and which plugins are *Uploaded*, so we know which are *Uploaded* though this specific term is never used in the Confluence UI. Instead, they are called *Dynamic*.

Upgrading plugins

- *Core* plugins cannot be upgraded.
- *WEB-INF/lib* plugins can be upgraded by replacing the JAR in WEB-INF/lib and restarting Confluence.
- *Bundled* plugins can be upgraded using the Plugin Manager or the Plugin Repository Client. A new plugin jar is uploaded and stored as a *Uploaded* plugin. Confluence compares the version number with the *Bundled* plugin and uses the newer.
- *Uploaded* plugins are upgradable using the Plugin Manager or the Plugin Repository Client. When a new plugin jar is uploaded, the previous version is discarded from the database and the `$CONFLUENCE_HOME/plugin-cache`.

RELATED TOPICS

Removing Malfunctioning Plugins

Removing Malfunctioning Plugins

Confluence goes to some lengths to prevent itself being unusable due to a problematic plugin. However, sometimes a plugin will manage to do this anyway.

Plugin Loading Strategies

1. Read through Plugin loading strategies in Confluence.
2. Determine where your plugin is loaded. The usual options are:
   a. The PLUGINDATA table on the database
   b. The <confluence-home>/bundled-plugins folder
   c. The <confluence-home>/plugin-cache folder
   d. The <confluence-home>/plugins-osgi-cache folder
   e. The <confluence-home>/plugins-temp folder
   f. The <confluence-install>/confluence/WEB-INF/lib folder (deprecated approach)

Check these locations when troubleshooting plugin loading issues.

> ✅ Check the How to display classpath utility for tips on what's loading, and the Knowledge Base Article on plugin malfunctioning.

Deleting a plugin from the Database

To remove a plugin from Confluence when Confluence is not running:

1. Connect to the Confluence database.
2. Run the following SQL statement in your database:

```
1.select plugindataid, pluginkey,filename,lastmoddate from plugindata;
```

3. After you have found the plugindataid for the offending plugin, please run the following:

```
1.delete from plugindata where plugindataid='XXXXXX';
```

where XXXXXX is the plugindataid value.

4. Restart Confluence.

Deleting a Bundled Plugin

Bundled plugins can be administered from the Plugins console from Administration >> Plugins. You can upload or disable them there.

Bundled plugins are included in a zip of jars called **atlassian-bundled-plugins.zip** which is on the main Confluence classpath, in a resources directory - `<confluence-install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup`. At Confluence startup, they are extracted and copied into the **$CONFLUENCE_HOME/bundled-plugins** directory, from whence they are loaded. To remove a bundled plugin (you shouldn't normally have to do this), remove the plugin from the atlassian-bundled-plugins.zip file and the bundled-plugins directory, otherwise Confluence will just put it back in place on the next startup. In versions later than 2.6, you'll have to recreate the .jar file (if the jar file is from the lib folder) or recreate the zip folder(if its in the classes folder). Bundled plugins can be upgraded or disabled.

If you need to remove a bundled plugin, check to see if you have duplicates in the `<confluence-home>/bundled-plugins` or `<confluence-home>/plugin-cache` directory.

Usually, the problem is that an old plugin is getting loaded along with the properly bundled one, but if you need to remove a bundled plugin, check Plugin loading strategies in Confluence.

Enabling and Configuring Macros

Macros allow you to perform programmatic functions within a page, and can be used for generating more complex content structures.

Generally speaking, a macro is simply a command wrapped inside curly braces {...}. To learn how to write your own macro, or use macros written by other people, read the Confluence Plugin Guide.

RELATED TOPICS:

- Configuring a URL Whitelist
- Configuring the userlister Macro
- Editing and Removing macros
- Enabling HTML macros
    - Enabling the html-include Macro
- Troubleshooting the Gallery Macro

Configuring a URL Whitelist

The RSS and HTML-include macros are used to include content dynamically from other websites onto a Confluence page. The included content may possibly be malicious or harmful to your Confluence instance.

Confluence administrators can set up a list of trusted URLs, thus limiting the locations from which the RSS macro and the HTML-include macro can draw their content.

The form below allows you to define specific URLs and/or URL patterns which are trusted, or to allow inclusion from all URLs without restriction.

To configure the URL whitelist,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Configure Whitelist**' in the left-hand panel.
3. The 'Configure Whitelist' screen will appear, as shown in the screenshot below.
4. Select one of the radio buttons as follows:
   - **Allow all domains** — There will be no restrictions to the content which can be included onto your Confluence pages.
   - **Restrict to listed domains** — Confluence will allow content from trusted URLs only. When you select this option, a textbox will open allowing you to enter specific URLs and/or URL patterns. Enter one or more URLs, each on its own line. You can enter the full URL, or use the pattern matching rules described below.
5. Click '**Save**'.

Screenshot: Configuring a URL whitelist

URL Pattern-Matching Rules

Enter one URL or URL pattern per line. You can enter a full URL or use pattern-matching as described below:

- If the rule starts with an equals sign (=), only the exact URL following the '=' will be allowed.
- If the rule starts with a slash (/) then the whole rule will be treated as a regular expression.
- Otherwise, any asterisk (*) will be treated as a wildcard to match one or more characters.

What Happens to a Page Containing a Disallowed URL?

A user can add the RSS macro or the HTML-include macro to a Confluence page. The macro code includes a URL from which the content is drawn. When the page is displayed, Confluence will check the URL against the whitelist. If the URL is not allowed, Confluence will display an error message on the page.

The error message says that Confluence "could not access the content at the URL because it is not from an allowed source" and displays the offending URL. If the person viewing the page is a Confluence Administrator, they will also see a link to the Administration page where they can configure the URL whitelist.

Here is an example of the error message, including the link shown only to Confluence Administrators:

Here is an example of the error message, but without the link.

Could not access the content at the URL because it is not from an allowed source.

http://ffeathers.wordpress.com

You may contact your site administrator and request that this URL be added to the list of allowed sources.

Notes

Some things to be aware of:

- By default, the RSS and HTML-include macros are disabled in Confluence. A System Administrator can enable them on the '**Plugins**' screen of the Confluence Administration Console.
- A user who has the 'Confluence Administrator' permission, but not necessarily the 'System Administrator' permission, can configure the URL whitelist (for the HTML-include and RSS macros).

RELATED TOPICS

Enabling HTML macros
RSS Feed Macro
HTML Include Macro

Administrators Guide Home   Confluence Documentation Home

Configuring the userlister Macro

The userlister macro has an optional 'online' parameter. If the '**User Listener**' plugin is configured to allow this feature, then the page author can specify 'online=true' to show a list of all online users.

You need to have System Administrator permissions in order to perform this function.

To enable the 'online' filter in the `userlister` macro,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will list the currently installed plugins.
3. Scroll down and click the '**User Listener**' link. The User Listener plugin panel will appear at the top of the screen.
4. Enable the '**User Log In Listener**' module by clicking the '**Enable**' link on its right.
5. Restart Confluence.

List of online users can be misleading
When the parameter 'online=true' is used, Confluence uses a context listener to generate the list of online users. A context listener is a J2EE term for something that listens for events in the application server. We listen for session open and close events, so a user is 'online' if they have a session on the application server. Some application servers don't correctly despatch close events for sessions – in these cases, the list of online users may be misleading.

Screenshot: Enabling the User Log In Listener

User Listener

Vendor: Atlassian Software Systems
Plugin Version: 2.1

A plugin which reports on Users, per group, within Confluence

□ Disable plugin

userlister Outputs lists of users, whether entirely or in specified groups	Disable
User Log in Listener Informs the UserLister macro when users log in or out of Confluence.	Enable

RELATED TOPICS

Userlister Macro
Enabling and Configuring Macros

Editing and Removing macros

ⓘ You need to have System Administrator permissions in order to perform this function.

To edit or remove a user macro,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**User Macros**' in the left-hand panel. This will list the currently configured user macros with options to '**Edit**' or '**Remove**' each macro.
    - Click '**Edit**'. This will display the edit screen for the macro. [Make changes] in the '**template**' input field and click '**Save**'.
      ⓘ Templates are in HTML, not wiki markup.

    - Click '**Remove**' to delete the macro.

RELATED TOPICS

📄 Enabling HTML macros

📄 Include Page Macro

📄 Editing and Removing macros

📄 Enabling the html-include Macro

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Enabling HTML macros

The {html} macro allows you to use HTML code within a Confluence page.

The {html-include} macro allows you to include the contents of an HTML file in a Confluence page.

CAUTION: Including unknown HTML inside a webpage is dangerous. Because HTML can contain active scripting components, it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.
By default, the HTML macros are disabled. You should only turn on these macros if you trust all your users not to attempt to exploit them.

ⓘ You need to have System Administrator permissions in order to perform this function.

To enable the HTML macros,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will display the installed plugins active for this Confluence installation.

3. Click' '**HTML macros**', then click '**Enable Plugin**'

RELATED TOPICS

Enabling HTML macros

Include Page Macro

Editing and Removing macros

Enabling the html-include Macro

Administrators Guide Home Confluence Documentation Home

Enabling the html-include Macro

The {html-include} macro allows you to include the content of an HTML file in a Confluence page. This page tells you how to enable the macro, so that it is available on your Confluence site. For help on using the macro, see HTML Include Macro.

> **CAUTION: Including unknown HTML inside a web page is dangerous.**
> Because HTML can contain active scripting components, it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.

Enabling the HTML Macros

By default, the HTML macros are disabled. You should only turn on these macros if you trust all your users not to attempt to exploit them.

 You need to have System Administrator permissions in order to perform this function.

To enable the HTML macros,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will display the installed plugins active for this Confluence installation.

3. Click' '**HTML macros**', then click '**Enable Plugin**'.

RELATED TOPICS

HTML Include Macro

Enabling HTML macros

Include Page Macro

Editing and Removing macros

Enabling the html-include Macro

Administrators Guide Home Confluence Documentation Home

Troubleshooting the Gallery Macro

Gallery Macro

The full list of parameters is shown in the following table.

Parameter	Default	Description
Gallery Title (`title`)	Nothing	Specify a title for your gallery.
Number of Columns (`columns`)	4	Specify the number of columns for your table.
Images to Exclude (`exclude`)	No exclusions i.e. include all the pictures on the page.	The gallery will ignore any pictures specified by `exclude=picture file name` i.e. they will not be included in the gallery. You can specify more than one picture, separated by commas. Example: `exclude=my picture.png,my picture2.gif`
Include these Images Only (`include`)	Include all the pictures on the page.	If you specifically include one or more pictures, the gallery will show only those pictures. Format is `include=picture file name`. You can specify more than one picture, separated by commas. Example: `include=my picture.png,my picture2.gif`
Use Images in these Pages (`page`)	If no page is specified, the gallery displays the images attached to the page containing the macro.	Specify the title of the page which contains the images you want displayed. If the page is in the same space as the page containing the macro, use the format `page=My Page Name`. To specify a page in a different space, use `page=SPACEKEY:My Page Name`, such as `page=DOC:Gallery Macro`
Reverse Sort (`reverseSort`)	Nothing, i.e. sort order is ascending	Used in conjunction with 'sort' parameter above. Use 'reverseSort' to reverse the sort order, from ascending to descending.
Sort Images By (`sort`)	None i.e. the sort order is unspecified and therefore unpredictable.	Specify an attribute to sort the images by that attribute. Sort order is ascending, unless you specify the 'reverseSort' parameter (see below). Options are: • '**name**' – file name. • '**comment**' – comment linked to the attached file. • '**date**' – date/time last modified. • '**size**' – size of the attached file.

> ℹ️ If the actual name of an attachment file or page contains a comma, you can refer to it in the **exclude**, **include**, or **page** parameters above by enclosing it in single or double quotes, for example `"this,that.jpg", theother.png`.

For more information, refer to Gallery Macro.

Troubleshooting

If you encounter the following error message: `System does not support thumbnails: no JDK image support` then ensure that you have following system property available for your JVM:

```
1.JAVA_OPTS=-Djava.awt.headless=true
```

Also see CONF-1737

> ⚠️ Please note that gallery-ext.jar is available at CONF-6620

Configuring the Office Connector

The Office Connector is a Confluence plugin that allows Confluence users to interact with Microsoft Office and Open Office in various ways. You can display content from Office documents on a wiki page, import content from an Office document into Confluence, and edit a Confluence page in Microsoft Word. Please refer to the User Guide for details of these interactions.

A System Administrator can enable or disable parts of the Office Connector and can configure options as described below.

On this page:

- Enabling and Disabling the Office Connector and its Modules
- Configuring the Office Connector Options

Enabling and Disabling the Office Connector and its Modules

A System Administrator can install, enable or disable plugins and plugin modules. You can read a general overview in Installing Plugins and Macros.

The Office Connector is bundled with Confluence 2.10 and later, so you should not need to install it. But you may wish to enable or disable some of its modules.

To enable or disable the Office Connector and its modules,

1. Select '**Plugins**' in the left-hand panel of the Confluence Administration Console.
2. Search the page for 'Office Connector plugin' and select the link.
3. The '**Office Connector plugin**' panel will appear near the top centre of the page, as shown in the screenshot below.
4. Now you can do one of the following:
   - **Configure plugin** — This will take you to the separate plugin configuration screen described below.
   - **Disable plugin** — Click this link if you want to disable all modules of the plugin, but leave the plugin installed on your Confluence site.
   - **Uninstall plugin** — Click this link if you want to remove the Office Connector permanently from your Confluence site. To restore it at a later date, you will need to re-install it from the Confluence Plugin Repository.
   - You can also enable or disable one or more of the Office Connector modules, as described in the table below.

Screenshot: Enabling the Office Connector plugin and its modules

The following modules are available for the Office Connector plugin:

Module	Description
OC Settings Manager	Component to read and write persistent settings for the Office Connector.
Slide Cache Manager	Component to cache slide-based conversions when displaying PowerPoint and PDF documents.
Html Cache Manager	Component to cache HTML-based conversions when displaying Word and Excel documents.
File Cache Cleanup Job	This module is a recurring task that cleans up the Office Connector file cache.
File Cache Cleanup	This module is the trigger for the File Cache Cleanup Job.
Office Connector administration link	This module supplies the 'Office Connector Configuration' link in the left-hand panel of the Confluence Administration Console. The link gives access to the plugin configuration screen described below.
Link for previewing a search result	This modules supplies the 'View' link which appears next to attachments displayed in search results, where the attachment is an Office document.
Link for previewing an attachment	This modules supplies the 'View' link which appears next to attachments displayed on the 'Attachments' view of a page, where the attachment is an Office document.
Edit in Word UI on page actions	This module supplies an 'Edit in Word' icon which appears in older versions of Confluence, at the top right of the page with other icons such as the 'Browse Space' and 'Add Page' icons. Not relevant to Confluence 2.10 or later, except for custom themes. You can configure the location of this option on the Office Connector configuration screen, as described below.
Edit in Word UI on drop down menu	This modules supplies the 'Edit in Word' link which appears in the Confluence 'Tools' dropdown menu.
editinwordlink	This modules supplies the {editinwordlink} macro. See Edit in Word Link Macro.
viewfile	This module supplies the {viewfile} macro. See View File Macro.
viewdoc	This module supplies the Word document component of the {viewfile} macro.
viewxls	This module supplies the Excel document component of the {viewfile} macro.
viewppt	This module supplies the PowerPoint document component of the {viewfile} macro.
viewpdf	This module supplies the PDF document component of the {viewfile} macro.
editgrid	This module is used to migrate editgrid users to the Office Connector.
Edit in Word UI on page tabs	This module supplies an 'Edit in Word' tab which appears in older versions of Confluence, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs. Not relevant to Confluence 2.10 or later, except for custom themes. You can configure the location of this option on the Office Connector configuration screen, as described below.
Import Word UI on page tabs	This module supplies a 'Doc Import' tab which appears in older versions of Confluence, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs. Not relevant to Confluence 2.10 or later, except for custom themes.
Import Word UI on drop down menu	This modules supplies the 'Doc Import' link which appears in the Confluence 'Tools' dropdown menu.
Edit in Office javascript resource	This module contains the javascript resources for launching the desktop applications for editing Office documents.
Office Connector Servlet	This module allows Confluence users to edit their Confluence pages in Microsoft Word. It performs the conversion to and from Word.
Office Authenticator Filter	This module authenticates HTTP requests from Office applications.
PPT slide web service	This module allows Confluence users to view a PowerPoint presentation on a wiki page. It provides the slide images to the Flash control which displays the slides on the wiki page.

DOC and XLS image cache web service	This module is required if Confluence users want to view a Word document or an Excel spreadsheet on a wiki page. It allows images to be stored in a cache on the server, so that they can be retrieved when the browser renders the HTML page.
Office Connector Actions	This module must be enabled if the Office Connector is used.

Configuring the Office Connector Options

A Confluence administrator can set the options described below, to determine the behaviour of the Office Connector on your Confluence site.

To set the configuration options for the Office Connector,

1. Select '**Office Connector Configuration**' in the left-hand panel of the Confluence Administration Console.
2. The '**Configure Office Connector plugin**' screen will appear. Set the configuration options as described in the table below.

Screenshot: Configuring the Office Connector options

The configuration options are described in the table below:

Option	Default Value	Description
Edit in word button location	'Tools' menu.	This setting determines the location of the '**Edit in Word**' option on the Confluence menus or screens. You may want to change the location if you are using a theme which does not support the default location. Note that you must ensure that the relevant plugin modules are enabled, as described above . Available settings are: • Page action icon — An icon at the top right of the page with other icons such as the 'Browse Space' and 'Add Page' icons. • View page tab — A page tab, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs.
Warnings: Show a warning before allowing a user to perform an import	Disabled	If this option is enabled, the user will receive a warning when importing a Word document. The warning will tell the user when they are about to overwrite existing content.

Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	If this option is enabled, a Confluence page created from an imported Word document will use the {footnote} macro from Adaptavist to render any footnotes contained in the document. Note that you will need to install the Footnotes Plugin onto your Confluence site. For more information about this plugin and macro, please refer to the Footnotes Plugin.
Authentication: Allow authentication tokens in the URL path	Disabled	If this option is enabled, the Office Connector will use authentication tokens in the URL.
Temporary storage for viewfile macro	The Confluence Home directory.	The {viewfile} macro will cache data temporarily. This option allows you to set the location of the cache. If you are running in a clustered environment, we recommend that you use the Cache in-memory setting.
Number of Conversion Queues	6	This is the maximum number of threads used to convert PowerPoint or PDF slide shows. You can use this setting to manage Confluence performance, by limiting the number of threads so that the Office Connector does not consume too many resources. Click the **Manage Queues** link to view attachments that are still pending conversion.

RELATED TOPICS

Office Connector Prerequisites
Office Connector Limitations and Known Issues
Working with the Office Connector
Installing Plugins and Macros

Operating Large or Mission-Critical Confluence Installations

This page gives guidelines for operational management teams who are responsible for a large Confluence installation, or for a Confluence installation which is crucial to the business of their organisation.

On this page:

- Introduction to this Page
    - Motivation for Presenting these Guidelines
    - Who should Read these Guidelines?
- Requirements of Large or Mission-Critical Confluence Installations
    - Dedicated Hardware for Confluence
    - Dedicated Qualified Staff
        - Operations Team with General Administrators
        - Network Staff
        - Database Staff
        - Developers
    - Constant Monitoring of Production Systems
    - Adherence to Strict Upgrade Procedures
    - Testing of Upgrades before Production Implementation
    - Load-Testing Environments
    - Tuning
        - Optimising your System
        - Limiting Third-Party Plugins
        - Selecting and Tuning your JVM
        - Customising Confluence to Optimise Performance

Introduction to this Page

Motivation for Presenting these Guidelines

Most Confluence installations start off small. Ten people in an early-adoption department use it for a couple of weeks. Everything works well and the good news starts spreading. Adoption increases throughout the organisation. More and more people use the wiki, and more and more rely on Confluence being up and running. After a while even the CEO starts blogging. And then a system outage occurs.

Now what?

Wikis like Confluence often grow into mission-critical applications within just a few months. Often adoption is so fast that IT departments haven't had the time to scale up their support.

We have assembled some requirements to help you make sure that your installation of Confluence can be mission critical. There are no surprises to be found here — all of the requirements would apply to any other piece of software that is mission critical within your organisation.

Who should Read these Guidelines?

The guidelines **do not apply to you** if you are using Confluence with just a few dozen users, and no one really minds if Confluence is down for a couple of hours because your database has crashed.

But if any one of the following applies to you, then these guidelines are **a must read for you**!

- The wiki has become your organisation's documentation base.
- Your users can't work properly when Confluence is down.
- Your boss or customer threatens to terminate your contract if you don't meet a strict service level agreement (SLA), such as 99.9% availability.

Requirements of Large or Mission-Critical Confluence Installations

Dedicated Hardware for Confluence

In a small work group with a few dozen or even hundreds of users, your Confluence installation can happily share the CPUs, memory and disks with other low-profile applications and a database.

But with thousands or even tens of thousands of users, you need dedicated hardware that runs Confluence and nothing else, and it needs to be fast hardware with plenty of RAM. While you *can* run Confluence in a virtualised environment such as VMware, we suggest you don't do it for mission-critical or high-load installations unless you are a real expert in virtualisation. Otherwise your other VMs might have performance problems which propagate to Confluence.

If you experience database-related problems, you should consider moving the Confluence database to a dedicated machine. Confluence itself can run queries that impact the performance of other applications, and other application problems or scheduled tasks can have an adverse affect on the usability of Confluence.

Dedicated Qualified Staff

If your Confluence installation is mission critical and your service level agreements require 24/7 up time, you need to be able to pinpoint problems quickly. You need qualified staff, dedicated to looking after Confluence, who are available during business hours and possibly beyond.

If you require assistance from the Atlassian Support team, you may need to answer some pretty technical questions to help us diagnose what is going on in your systems. Also keep in mind that Atlassian support assists you in finding problems in Confluence, but we can't help you administer your systems.

In particular, we recommend that you have dedicated staff in the roles listed below.

Operations Team with General Administrators

If your organisation relies on Confluence being up and running around the clock with very little downtime, you need people who can set up, maintain, tune and improve your Confluence installation. This requires at least one person, but ideally you will have a team of operational engineers.

If your wiki is mission critical, chances are that other IT systems within your organisation have already made it necessary to have such an operations team. So you will probably not need to hire someone specifically to administrate Confluence. But it is vital that supporting and maintaining Confluence is added to the list of responsibilities of that operations teams, and that you can get them to troubleshoot and analyse Confluence at short notice.

If problems arise and you need to contact Atlassian Support, these engineers will be our first point of contact. We may ask them to provide details of log files, application-server settings, monitoring systems, and so on.

Network Staff

If Confluence is mission critical for large numbers of users, it is vital that you have dedicated network staff available to track down problems when they arise.

A mission-critical installation will usually be used by hundreds or even thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone has made an undocumented change to the network and you don't have an expert around who can figure it out.

Again, this only applies to mission-critical systems. If you use Confluence for less critical collaboration and knowledge sharing, and a broken network cable causing a day's downtime is no major catastrophe, then you will not need dedicated networking staff.

Database Staff

If Confluence is mission critical for a large number of users, you need an experienced database administrator (DBA) available to troubleshoot database performance issues and other potential problems. It is dangerous not to have an experienced full-time DBA at hand at short notice when running a mission critical application. While small installations of Confluence basically work 'out of the box', any system that involves high load or high-availability requirements needs continual monitoring, optimising and fine tuning of the Confluence database. Database monitoring is no trivial task — it's not something that anyone can learn quickly.

Developers

You may have decided to customise Confluence by changing its source-code, or by writing your own plugins. If your server is mission-critical, you must nominate staff who will be responsible for that code, and they must be up for the task. Otherwise you might end up in a situation in which your server experiences downtimes because of custom code is broken, or does not work with a newer version of Confluence anymore, but you can't fix the problem because no one knows how the customized code works, and you can't uninstall it either because it has become critical for your Confluence usage pattern. Keep good track of changes, and have someone available to jump into action if there is a problem Don't let the summer intern write mission-critical plugins, unless you have more senior staff to maintain that code as long as it is in use.

Constant Monitoring of Production Systems

You will need to monitor your production systems constantly.

When the wiki is the lifeblood of your organisation, you need know exactly what is going on inside, so that you can plan for future needs and analyse potential bottlenecks.

Monitoring involves a number of essential tasks, including those listed below:

- Monitoring log files.
- Checking for HTTP-availability and performance (e.g. by getting the same page every five minutes and displaying the time on a graph).
- Looking at many different parameters such as load, connections, IO, database-trends, and so on.
- Charting long-term trends.
- Keeping an access log of requests to the web server. This is vital, especially when requesting performance-related support from Atlassian.

Monitoring a web application like Confluence implies also monitoring the subsystems it uses. Many outages and downtimes are caused by broken mail servers, databases running out of space, file systems filling up and so on. It is often possible to detect these trends way before the actual web application breaks down. Keep an eye on the file system, and if you see it is getting closer to 90% utilisation, you can mend the situation without Confluence breaking down. Or even if the worst case happens (e.g. the database breaks down and Confluence is affected straight away) then having the proper monitoring for the database server makes troubleshooting a lot easier.

> **Tools for Monitoring Confluence**
> At Atlassian we use Hyperic. But the list of monitoring systems is long and we can't recommend a specific product over the other. If your organisation has a monitoring system already, make sure you hook up Confluence to it. If you don't have a monitoring system yet, you need to install one as soon as you feel Confluence is mission critical.

As an example of what our monitoring UI looks like, have a look at this screenshot:
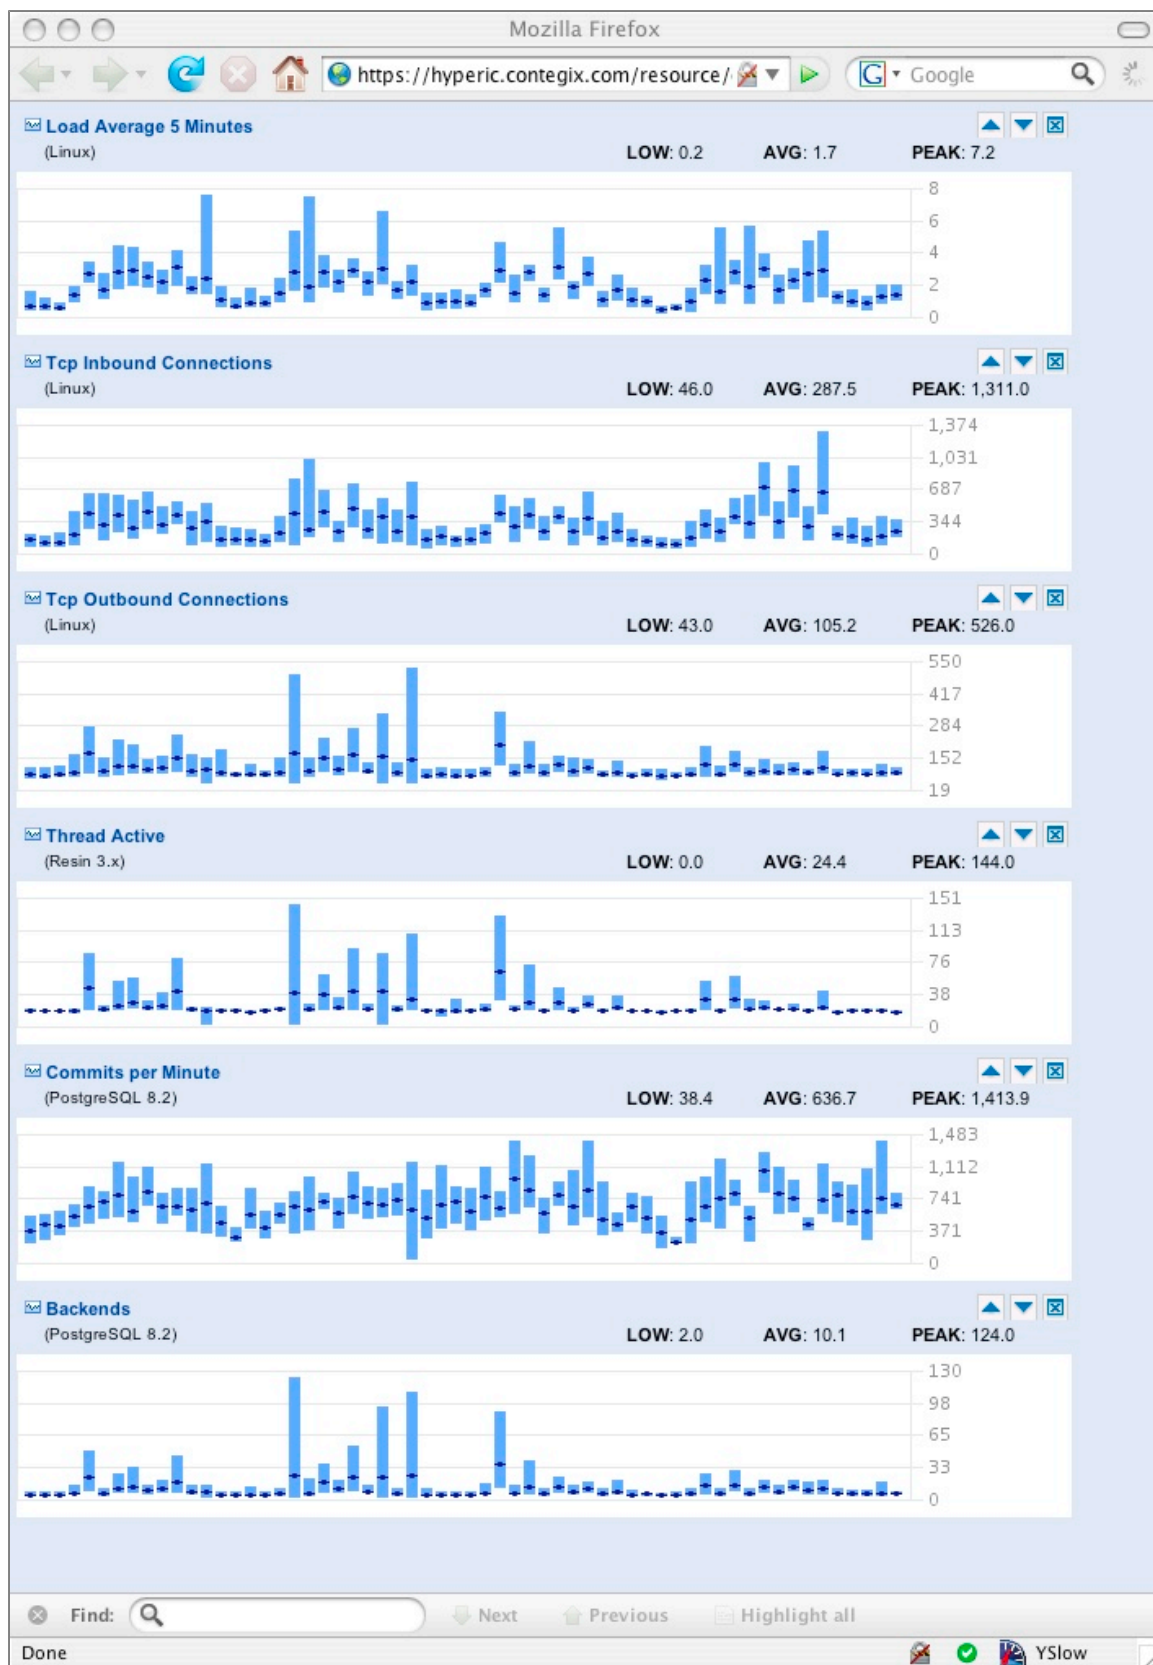
The following screenshot shows one of our sensors looking at the HTTP response times of our documentation wiki over the last 8 days. You can clearly see an incident four days ago. Having the graph (and regularly looking at it) allowed us to pinpoint the problem. We analysed the access logs and found that webpage-profiling had been enabled but not disabled again, which caused performance problems.
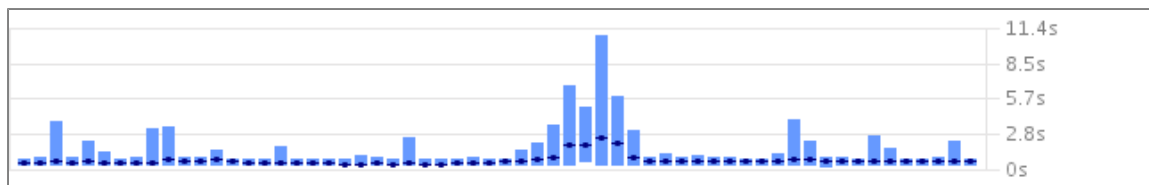
This page would get too long if we described all our monitoring sensors - but just to give you an impression, this is what we monitor on the JVM level alone.

JVM basics

- Current Loaded Classes
- Daemon Thread Count
- Heap Memory Committed
- Heap Memory Max
- Heap Memory Used
- Loaded Classes
- Loaded Classes per Minute
- Object Pending Finalization Count
- Peak Thread Count
- Thread Count
- Unloaded Classes
- Unloaded Classes per Minute

JVM garbage collection

- Collection Count
- Collection Count per Minute
- Collection Time
- Collection Time per Minute

JVM memory: (Metrics for Eden space, Old Gen, Survivor space, Perm Gen)

- Commited Memory
- Used Memory

We get the same level of detail for our database, for the file system, for the CPU, for the network, and so on. Not all of this is needed all the time. But if your company depends on an application, then the more information you have at your fingertips the better. Fortunately these metrics can be extracted quite easily once you have a monitoring system in place.

Adherence to Strict Upgrade Procedures

Your organisation will have its own upgrading procedure. Here are a few recommendations that you should add to your list:

- Our main recommendation: Never change more than one component at a time. Sometimes it may be tempting to upgrade the server hardware when you upgrade Confluence, but we recommend you don't do that. It makes pinpointing errors much more difficult. So, for example, don't upgrade hard disks in conjunction with a Confluence version upgrade, don't change the Confluence configuration at the same time as you upgrade your Apache software, and don't upgrade a major third-party plugin the day you move your database system to a new machine. The list is endless, these were just a few examples to get you thinking.
- After each upgrade step, run Confluence for a couple of days to check that everything is still fine.
- Keep track diligently of what you change, and when. It will be nearly impossible for us to help you if you can't tell us what exactly you changed at what time.
- Keep a copy of all log files produced during the upgrade, together with notes about what changed between successive restarts.

Always take careful note of the upgrade notes published with the Release Notes of each Confluence version, as well as the Confluence Upgrade Guide.

Example

Here you can see an extract of our change log for `http://confluence.atlassian.com` — the server that hosts this very page.

Sydney time	Server time	Event	Reason/Purpose (including JIRA issues)
	2008-03-25 22:18	Started upgrade to 2.8-m9-r3 (build #1314)	
	2008-03-25 22:25	App server brought down due to failed database upgrade	
	2008-03-26 00:51	Server brought back up after database restored from backup. Running 2.8-m9-r3.	
	2008-03-28 04:18	GC algorithm changed from concurrent to parallel collector. Max heap increased from 1.4 GB to 2.0 GB	

	2008-04-24	Hyperic agent started with connection to Resin.	
	2008-05-08 20:30 - 22:30	Manual updates to menu.css, comments.js and comments.css in webapp	Temporary fix for , which was impacting performance
	2008-05-12	Updated cache sizes for five caches, bounced server.	Cache efficiency was low on these caches.
2008-05-13 18:00-18:20	2008-05-13 03:00-03:20	Upgrade from Resin 3.0 to Tomcat 5.5	
2008-05-14 16:30-17:00		Upgrade from Confluence 2.8.1-rc2 to 2.8.1-rc3	
	2008-05-14 20:30	Install new cronjob as j2ee for automating access log analysis	

Testing of Upgrades before Production Implementation

You should test upgrades in a staging environment.

Before rolling out a new version of Confluence (or of the software or hardware that it uses, e.g. database systems, application servers, data storage), make sure that you test the upgrade with real data (e.g. a database dump) on a completely independent machine.

Here's an example of what such a test would pick up: The new release of Confluence may not be compatible with a custom third party plugin you have previously installed, thus breaking the plugin's functionality. You may not even know that anyone installed that plugin — but maybe many people are already using it. You'll want to find out about this before you actually roll out the new version of Confluence.

Here is an outline for a simple upgrade test:

1. Create a clone of your production environment, using a database dump to obtain a copy of the Confluence data. We'll call this your 'staging environment'.
2. Upgrade the staging environment to the new version of Confluence.
3. Ask a few selected users from different departments to check the pages they commonly access, but have them do it in the staging environment.

Hint: In addition to finding weirdnesses with plugins, this may also show whether training for new functionality is needed in some of the departments. The IT department staff may be able to handle the upgrade to a new version of Confluence without training, but perhaps the sales representatives who use the wiki less often will need some training.

> **Getting a license for your staging environment**
>
> > **Only a technical contact for your commercial/academic license is able to create a Developer license**
>
> Atlassian supplies 'developer' licenses which can be used by existing commercial license holders who wish to deploy non-production installations of our software to use in QA/staging environments. Developer licenses are free of charge to commercial license holders and, like our commercial offerings, they include 12 months of updates starting from the date of purchase of the commercial license.
>
> If you hold a commercial license, you can obtain a free developer license by performing the following:
>
> 1. Log in to your Atlassian account.
> 2. Under the "Licenses" heading, all of your licenses will be displayed. Click the plus sign next to a license to view its details.
> 3. Click the **'View Developer License'** link in the bottom right corner of the license detail panel, below your commercial license key.

Load-Testing Environments

Many customers ask us,

> *So, how many users and spaces can I put into Confluence, and what is the best hardware do to so?*

The answer is, 'It depends'.

It depends a lot on your use case. Confluence is so successful because it can cover a huge range of use cases. If most of your users only access Confluence infrequently, it is no problem to have 70 000 to 100 000 users. But if each user is a power-user who uses the system the whole day, there's a substantial decrease in number Confluence can take without tuning. If your pages are short, simple, and don't contain a lot of macros, then the situation will be vastly different from a system that relies heavily on macros, background-tasks, or other features.

If your system is large (for example serving more than 10 000 users or storing more than 1000 spaces) or mission-critical (which it could be

with as few as 1000 users who use it all the time) you need one or more more load-testing environments.

Even if your system is working nicely for 20 000 users right now, it might take just another 2000 users to push it over the edge.

We recommend the following basic procedure:

- Set up an environment that closely resembles your production environment.
- Gather statistics from your production system.
- Regularly apply a similar kind of load (and slightly higher) to the load-testing environment.
- Analyse how well Confluence scales for your usage patterns.

The Confluence development team has load-testing scripts available which you can use to simulate load. You can also contact Atlassian Support for more details.

Tuning

You may need to be able to tune your installation in the ways mentioned below.

Optimising your System

If you have large numbers of users, then downloading all the static content (CSS, default images, JavaScript-files) may result in a high additional load on the application server that can be offloaded to a caching web server.

Please refer to the following additional information:

- Our general Performance Tuning page.
- Information on configuring a large Confluence installation.

Limiting Third-Party Plugins

You may have to restrict the number of third-party plugins installed on your Confluence instance.

Most third-party plugins are not specifically written for high-load environments. What works fine in low-load environments could have unexpected and adverse effects when thousands of users are competing for your application server's CPU time or for database IO.

A common source of problems is access to database connections. If you have fewer users than database connections, it does not matter if an operation holds on to a database connection for two seconds while it downloads some data from the internet. With hundreds of concurrent users, this could quickly become a bottleneck.

Confluence itself is tested and optimised to handle high loads and avoids these kinds of problems. But if you install a number of plugins that have not been tested against high load, your system may become unstable.

We recommend that you load test the common use cases of each unofficial third-party plugin if your Confluence installation is mission critical. Only activate plugins that are vital to your business, and never allow experimental plugins onto your production system until they have been tested in a staging environment.

Selecting and Tuning your JVM

You should select your JVM carefully and you may need to be able to tune it.

The selection of the JVM for your large Confluence instance can have a huge impact on the performance perceived by the users. Between versions 1.4 and 6 of the Sun Java JVM there have been some impressive improvements in performance, especially under high concurrent load.

Here are some essential guidelines:

- Always run the most recent point release of your selected JVM.
- Where ever possible run the most recent major release from your selected JVM manufacturer. The Sun JVM version 6 is much faster than 1.4, especially under high loads.
- Tune your garbage collection algorithms. Experiment with different algorithms and settings to get the response times you desire in your environment. Here are some specific guidelines for Sun JVM in the Sun documentation:
    - Java 6
    - Java 5
    - Java 1.4

Customising Confluence to Optimise Performance

You may need to customise Confluence for performance reasons. Depending on your usage scenario, there may be ways to enhance Confluence performance that become necessary when you reach a certain level of usage.

Here are some things you might decide to do:

- Remove the display of the space list on the Dashboard. See Customising the dashboard for Administrators.
- Configure any search appliances or other crawlers which are configured to index the Confluence site:
    - These should be suitably rate limited.
    - Configure them to crawl only pages in the `/display/` URL path, and only current versions of pages.

Please refer to our general Performance Tuning page for more details.

RELATED TOPICS

Performance Tuning
Configuring a Large Confluence Installation
Confluence Clustering Overview
Requesting Performance Support
Administrators Guide
Configuration Guide
Server Hardware Requirements Guide
Managing Application Server Memory Settings

Performance Tuning

- Description
- Use the latest version of your tools
- Avoid swapping due to not enough RAM
- Careful about those other systems using the same infrastructure
- Choice of Database
- Database Connection Pool
- Database in general
- Database indexes
- Database Statistics and Query Analysers
- Cache Tuning
- Antivirus Software
- Enabling HTTP Compression
- Virtual Operating Systems
- Performance Testing
- Access logs
- Built-in Profiler
- Adjust Application Server Memory Settings
- Use A Web Server
- Parallel GC
- Troubleshoot possible memory leaks
- Some 3rd-party plugins were not written to scale to large enterprises' needs
    - RELATED TOPICS

> ℹ This document describes tuning your application for improved performance. It is not a guide for troubleshooting Confluence outages. Check Troubleshooting Confluence Hanging or Crashing for help if Confluence is crashing. NEW: Garbage Collector Performance Issues

Description

Like any server application, Confluence may require some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read Requesting Performance Support.

Use the latest version of your tools

Use the latest versions of your application servers and Java runtime environments. Newer versions are usually better optimized for performance. As an example, our internal performance tests show a **20% speed-up** (when viewing pages under load) between Tomcat 6 on Java 6 vs Tomcat 5.5 on Java 5 **out of the box**.

Avoid swapping due to not enough RAM

Always watch the swapping activity of your server. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection considerably and affect Confluence's performance. In clustered installations, swapping can lead to a Cluster Panic due to Performance Problems. This is because swapping causes the JVM to pause during Garbage Collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Careful about those other systems using the same infrastructure

It may sound tempting: Just have one powerful server hosting your database and/or application server, and run **all** your crucial programs on that server. **If** the system is set up perfectly, then you might be fine. Chances are however that you are missing something, and then one application's bug might start affecting other applications. So if Confluence is slow every day around noon, then maybe this is because

another application is using the shared database to generate complicated reports at that time? Either make sure applications can't harm each other despite sharing the same infrastructure, or get these systems untangled, for example by moving them to separate instances that can be controlled better.

Choice of Database

The **embedded database** that is provided with Confluence is meant only to be used for evaluation, not for production Confluence sites. After the evaluation finishes, you will certainly need to switch to an external relational database management system. Beyond this, we do not recommend any particular RDBMS over another. We recommend using what you are familiar with, because your ability to maintain the database will probably make far more difference to what you get out of it than the choice of database itself.

Database Connection Pool

If load on Confluence is high, you may need more simultaneous connections to the database.

- If you are using JNDI data-sources, you will do this in your application server's configuration files.
- If you have configured Confluence to access the database directly, you will need to manually edit the hibernate.c3p0.max_size property in the confluence.cfg.xml file in your confluence.home directory. After you have changed the URL in this file, restart Confluence.

To assess whether you need to tune your database connection pool, take thread dumps during different times (including peak usage). Inspect how many threads have concurrent database connections.

Database in general

If Confluence is running slowly, one of the most likely cause is that there is some kind of bottleneck in (or around) the database.

The first item you should check is the **"Database Latency"** field in the System Information tab in the admin console.

Database Connection Transaction Isolation	Read committed
Database Latency	0
	Confluence Usage

The latency is calculated by sending a trivial request to the database, querying a a table which is known to have only one column and one row. ("select * from CLUSTERSAFETY"). Obviously this query should be blazing fast, and return within 1 or 2 milliseconds. If the value displayed is between 3 and 5 milliseconds, you might already have an issue. If the value is above 10ms, then you **definitely** need to investigate and improve something! A few milliseconds may not sound so bad, but consider that Confluence sends quite a few database queries per page request, and those queries are a lot more complex too! High latency might stem from all sorts of problems (slow network, slow database, connection-pool contention, etc), so it's up to you to investigate. Don't stop improving until latency is below 2ms on average.

Obviously, latency is just the very first thing to look at. You may get zero latency and still have massive database problems, e.g. if your tables are poorly indexed. **So don't let a low latency fool you either.**

Database indexes

Especially if you have more than a few thousand active users, and all most obvious measures have been tried out but the database still seems to be under high load, you should consider enaging a database administrator (DBA) to tune the database specifically to the demands that your particular Confluence installation is placing on it. If you do not have a full-time DBA and can't even get one for temporary consulting, you may want to consult the database indexing advice that we have been gathering from customer reports and our own experience running and developing Confluence. The instructions on that page are for Oracle, but most of the indexes can be applied to (and will help with) any database.

(These database indexes are now created automatically when Confluence is installed, but existing installations upgrading to a more recent version may still need to add them manually)

Database Statistics and Query Analysers

Modern databases have query optimisers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimiser is performing. If the cost estimate is wildly inaccurate then you will need to run statistics collection on the database. The exact command will depend on your database and version. In most cases you can run statistics collection while Confluence is running, but due to the increased load on the database it's best to do this after normal hours or on a week-end.

Cache Tuning

To reduce the load on the database, and speed up many operations, Confluence keeps its own cache of data. Tuning the size of this cache may speed up Confluence (if the caches are too small), or reduce memory (if the caches are too big). Future versions of Confluence will allow you to tune the size of this cache from within the web application. Vote for tuning the cache from the UI and getting cache recommendations to encourage Atlassian to build this feature into a Confluence release.

Please have a look at our documentation on Cache Performance Tuning for information on how to tune Confluence caches.

Antivirus Software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

Enabling HTTP Compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider enabling HTTP compression. This may also be useful when running an external facing instance to reduce your bandwidth costs.

⚠ Take note of the known issues with HTTP compression in versions of Confluence prior to 2.8, which may result in high memory consumption.

Virtual Operating Systems

Virtual Environments such as VMWare can cause Confluence CPU to spike. Run Confluence on a native OS. See our List Of Supported Operating Systems.

Performance Testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behaviour. Learn about how to test performance issues using the Performance Testing Scripts.

Access logs

You can find out which pages are slow and which users are accessing them by enabling Confluence's built-in access logging.

Built-in Profiler

You can identify the cause of page delays using Confluence's built-in profiler according to Troubleshooting Slow Performance Using Page Request Profiling.

Adjust Application Server Memory Settings

See Managing Application Server Memory Settings.

Use A Web Server

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to Running Confluence behind Apache.

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

Parallel GC

If you have multiple CPU's on your server, you can add -XX:+UseParallelOldGC to your JAVA_OPTS options. This will allow garbage collection of the Tenured Space to happen in parallel with the application and can boost performance and can reduce slow performance spikes. For more information, please refer to our detailed page on Garbage Collector Performance Issues, and Sun's summary of collectors.

Troubleshoot possible memory leaks

Some external plugins, usually ones that have been written a long time ago and that are not actively maintained anymore, have been reported to consume memory and never return it. Ultimately this can lead to a crash, but first this manifests as reduced performance. The Troubleshooting Confluence Hanging or Crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang.

Some 3rd-party plugins were not written to scale to large enterprises' needs

Confluence has been optimized to work under high load and with many pages. Some 3rd party plugins however have been written with small size companies in mind, and can't cope with large numbers of concurrent users, or large numbers of pages and permissions, or large numbers of spaces. It is impossible to tell which ones will fail under which conditions, but it will always help to turn off 3rd-party plugins that are not strictly mission-critical while investigating performance issues.

RELATED TOPICS

Garbage Collector Performance Issues
Cache Performance Tuning
Cache Performance Tuning for Specific Problems
Performance Testing Scripts
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview
Requesting Performance Support
Administrators Guide
Configuration Guide

Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches. It is essential for the administrator of a large production installation of Confluence to tune the caches to suit its environment. There are several configurable parameters for each of the cache regions, most notably cache size, cache expiry delay and eviction policy. In the majority of the cases, cache size is the parameter you would want to change. Fortunately, from Confluence 3.0, it is very easy to adjust cache sizes through the Administration Console. However, if you need to modify parameters other than a cache size, you would need to modify the relevant configuration files manually.

> ℹ️ If you only need to modify Confluence's maximum cache sizes, you can do this through the Cache Statistics feature of the Administration Console.

The cache performance information for your Confluence installation is available under **Administration > Cache Statistics**. More information about the numbers displayed here is available on Cache Statistics.

On this page:

- Cache tuning example
- Finding the configuration file
- Cache Key Mappings
- Standard Editions of Confluence
    - Understanding the Ehcache Configuration File
    - Converting your Coherence configuration to Ehcache
- Clustered Editions of Confluence
    - Understanding the Coherence configuration file
    - Defining Caching Scheme Mappings in Coherence Cache config file
- Important Caches
- Cache Tuning Follow-Up

Cache tuning example

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Used	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279

The caches above are of size 1000 (meaning that it can contain up to 1000 objects), which is the default size for caches in the default cache scheme. Refer to Confluence Cache Schemes for more explanation.

You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be eliminated from the cache. Cache expiry may be based on time or on frequency of use.

ⓘ There is not much that you can do with a cache that has both a low percentage of usage and effectiveness. Over time, as the cache is populated with more objects and repeat requests for them are made, the cache's effectiveness will increase.

Finding the configuration file

The caches are configured in `ehcache.xml` (for standard editions) or `confluence-coherence-cache-config-clustered.xml` (for clustered editions) which is stored in `<confluence-home>/config/`.

> ⓘ **Oracle Coherence Licensing Change:**
>
> - Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
>   - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
>
>   ⚠ **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
>   - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
> - For more information about these changes, please refer to the Coherence License Changes document.
> - If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later after late September 2009, please ensure that you download only a clustered edition of Confluence and please refer to the Confluence 3.0.1 Upgrade Notes for additional upgrade information.

Cache Key Mappings

The cache configuration file configures caches by their keys. When you move your mouse over the the cache names displayed on the cache statistics page, a tooltip will indicate the actual cache key for that cache name.
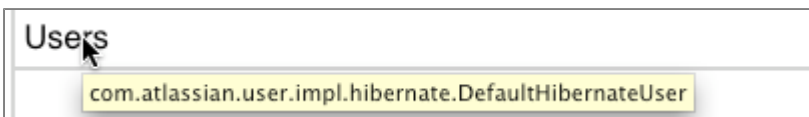
Using our example from the table above, if we were to modify parameters for the `Users` cache we would need to change the cache with the key `com.atlassian.user.impl.hibernate.DefaultHibernateUser`. Do not get confused with `Users (External Mappings)` and `Users (External Groups)` which are in themselves, two separate caches. "Users" is the friendly name for `com.atlassian.user.impl.hibernate.DefaultHibernateUser`.

Standard Editions of Confluence

In standard editions of Confluence, the caching layer is Ehcache.

Understanding the Ehcache Configuration File

For more information about the Ehcache configuration file and a full reference on Ehcache configuration, please refer to the Ehcache configuration documentation.

Converting your Coherence configuration to Ehcache

> ⚠ This section only applies to customers who:
>
> - Have an installation of Confluence that was downloaded before the 4[th] of September 2009.
> - Intend to (or have already) upgraded to Confluence 3.0.1 or later (or to Confluence versions 2.6.3, 2.7.4, 2.8.3, 2.9.3 and 2.10.4).
> - Will use a non-clustered Confluence license for the Confluence upgrade.
> - Have implemented customisations to their Confluence installation's cache configuration file (`confluence-coherence-cache-config.xml`).

To maintain your existing cache configuration file settings, you will need to transfer any cache customisations you have implemented in the Coherence cache configuration file (`confluence-coherence-cache-config.xml`) to the relevant entries in the Ehcache cache configuration file (`ehcache.xml`).

Each cache has a `cache-mapping` element in the Coherence file (of which there is an equivalent `cache` element in the `ehcache.xml` file). Unfortunately, copying across your customisations is not quite a straightforward process because the Coherence file defines several 'caching schemes' to store the actual cache values, which in turn are referenced by the `cache-mapping` elements. In contrast, the `ehcache.xml` file does not support caching schemes and a cache's values are expressed explicitly in separate parameters of a `cache` element.

To convert your Coherence cache configuration file customisations across to the equivalent Ehcache file:

1. Open both the `confluence-coherence-cache-config.xml` and `ehcache.xml` files in a text editor. These files are located in the `<confluence-home>/config` directory.

   If you implemented your customisations in a version of Confluence prior to 3.0, you will most likely find the `confluence-coherence-cache-config.xml` file in the `<confluence-install>/confluence/WEB-INF/classes` directory.

2. In the customised `confluence-coherence-cache-config.xml` file:

   a. Identify the caching schemes that were customised in this file and make a note of the values of all its child elements.

      Typically, each caching scheme is located inside a `local-scheme` element and all of these are enclosed within the `cache-schemes` element, which appears towards the end of this file.

   b. Note each customised caching scheme by the content of its `scheme-name` element.

   c. For each `cache-mapping` element (which typically appears towards the top of this file), identify if it has a `scheme-name` element whose content matches one noted in the previous step and if so, make a note of its associated `cache-name` element.

3. In the `ehcache.xml` file:

   a. Identify each `cache` element whose 'name' parameter matches the `cache-name` elements noted in step '2c'.

   b. Using the mappings table below, apply the values noted in step '2a' to the appropriate parameters of the `cache` elements identified in the previous step ('3a').

Mappings table showing how elements of the Coherence cache configuration file map to parameters of the equivalent Ehcache file.

Coherence Element	Ehcache Attribute
`high-units`	`maxElementsInMemory`
`expiry-delay > 0s`	`timeToIdleSeconds` - Use this attribute for expiry delays greater than 0s along with the `eternal` attribute set to 'false'
`expiry-delay = 0s`	`eternal` - For expiry delays of 0s, set this attribute to 'true'.

Clustered Editions of Confluence

Understanding the Coherence configuration file

The Coherence configuration file is a mapping of *cache keys* to *cache schemes*. Each cache scheme controls the expiry, eviction policy and size of the caches linked to it. A cache scheme can extend another scheme.

For a full reference, see the Oracle's Coherence cache configuration documentation.

Defining Caching Scheme Mappings in Coherence Cache config file

If a cache key does not have an explicit definition in the caching scheme mappings (defined in `confluence-coherence-cache-config.xml`) then it will use the "default" `cache-mapping`.

In our example, `com.atlassian.user.impl.hibernate.DefaultHibernateUser` is not explicitly defined in the caching scheme mappings. Hence to increase the expiry-delay to 2 hours, we will need to define the mapping ourselves and add the following within the `<caching-scheme-mapping>...</caching-scheme-mapping>` tags:

```
1.<cache-mapping>
2.<cache-name>com.atlassian.user.impl.hibernate.DefaultHibernateUser</cache-name>
3.<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
4.</cache-mapping>
```

Then we will need to define a cache schema with name `cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser` within `<caching-schemes>...</caching-schemes>` tags.

```
1.<local-scheme>
2.<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
3.<scheme-ref>default</scheme-ref>
4.<high-units>10000</high-units>
5.<expiry-delay>7200</expiry-delay>
6.</local-scheme>
```

It's possible to define a local-scheme mapping for a cache key without defining certain parameters (e.g. `<high-units>`). In such a cases, their parameters will be inherited from `scheme-ref` scheme, which is the `default` scheme in our case.

Important Caches

⚠ The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and Percent Used categories in the cache for more specific assessments.

- **com.atlassian.confluence.core.ContentEntityObject** (known as **Content Objects** cache)
  should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **com.atlassian.confluence.core.ContentEntityObject.bodyContents** (known as **Content Body Mappings** cache) should be set to at least 20% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **com.atlassian.confluence.security.PermissionCheckDispatcher.isPermitted()** (known as **User Authorized URLs** cache)
  should be set to at least the number of concurrent users you expect to access Confluence at the same time
- **com.atlassian.user.impl.hibernate.DefaultHibernateUser** (known as **Users** cache)
  should be set to the number of users you have: `select count (*) from users`. Note that by default, this will also control the LDAP user's cache, including expiration.
- **com.atlassian.confluence.security.SpacePermission** (known as **Permissions** cache)
  should be set to the number of space permissions in your deployment (a good rule of thumb is 20 times the number of spaces). You can find the number of space permissions using the query `select count(*) from SPACEPERMISSIONS`.

Cache Tuning Follow-Up

After you have made changes to your cache config, doing a follow up on the changes in the next week or after the expected performance spike would be important.

Make sure that you take a screenshot of the cache statistics before and after the change. Then compare them with the cache statistics in the later period where performance improvement is expected.

✅ You can monitor what's in the cache by using a JSP included in the Confluence distribution. Browse to <base-URL>/admin/cachecontents.jsp to monitor the cache contents.

RELATED TOPICS

Cache Performance Tuning for Specific Problems
Confluence Cache Schemes
Performance Testing Scripts
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview
Requesting Performance Support
Administrators Guide
Configuration Guide

Cache Performance Tuning for Specific Problems

The following are more specific performance problems that can be resolved from tuning the cache.

LDAP cache sizes and expiry does not appear to be picked up.

This is a known problem, please refer to CONF-11858 for the solution.

"Edit Page" screen takes a long time to load

If your installation of Confluence is suffering from this problem, it may be due to a insufficient SpacePermissions cache size. To address this problem, first determine the number of space permission objects in your Confluence instance. You can do this by running this query against your database:

```
> select count(*) from SPACEPERMISSIONS
```

Now locate the cache entry for SpacePermissions in your `confluence-coherence-cache-config.xml`:

```
1.<local-scheme>
2.
<scheme-name>cache:com.atlassian.confluence.security.CachingSpacePermissionManager.permissions</schem
3.                    <scheme-ref>default</scheme-ref>
4.                    <high-units>10000</high-units>
5.                    <expiry-delay>0s</expiry-delay>
6.</local-scheme>
```

Adjust the **maxElementsInMemory** or **high-units** property to the number of space permissions you have (in the example above, I've used 10000). Also, just as important, you need to adjust the **timeToLiveSeconds** or **expiry-delay** property to `0`.

Note: 10K of space permissions consumes approximately 8MB of memory. Please ensure there is enough memory allocated to your instance to cater for this.

How to set specific cache settings

1. Find the cache name from the cache name mappings:
   - For **Confluence 2.5.x and earlier**, the cache name mappings are in file `confluence/WEB-INF/classes/com/atlassian/confluence/admin/actions/cache-name-mappings.proper`.

   - For **Confluence 2.6.0 and later**, you will find the cache name mappings in the file `com/atlassian/confluence/core/ConfluenceActionSupport.properties` which is packed into the `confluence-2.x.*.jar file`.
2. Find the appropriate `<cache-mapping>` tag in `confluence-coherence-cache-config.xml` or `confluence-coherence-cache-config-clustered.xml`. If the tag doesn't exist, you can create it within the `<caching-scheme-mapping>` tag.

   > ⚠ Attached to this page are corrected copies of confluence-coherence-cache-config.xml and confluence-coherence-cache-config-clustered.xml. These are updated from a bug CONF-11857.

3. The `<scheme-name>` will correspond to a `<local-scheme>` tag below. It refers to a scheme reference. Either change the high-units tag in the scheme reference, or add a high-units tag to override the scheme reference. For example, the following tag would change the Content Bodies cache from the default 1000 units to 2000 units:

   ```
   <local-scheme>
   <scheme-name>cache:com.atlassian.confluence.core.ContentEntityObject.bodyContents</scheme-name
   ```

   Another popular cache to change is the LDAP related User cache:

   ```
   <local-scheme>
   <scheme-name>user</scheme-name>
   <scheme-ref>default</scheme-ref>
   <high-units>5000</high-units>
   <expiry-delay>300s</expiry-delay>
   </local-scheme>
   ```

4. After updating the appropriate file, you do not need to repack it into the jar to use it. You can simply place the file in your `confluence/WEB-INF/classes/` directory. The file in this directory will override the settings in your jar file. If you want to back out the changes, you only need to remove the file from your `confluence/WEB-INF/classes/` directory — then the default values in the `confluence-coherence-cache-config.xml` located in your jar file will apply.

You can find more information about configuring the Coherence cache in the Coherence cache documentation.

RELATED TOPICS

Cache Performance Tuning
Performance Testing Scripts
Confluence Cache Schemes
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview
Requesting Performance Support
Administrators Guide
Configuration Guide

Confluence Cache Schemes

Default Scheme

If a cache has not been defined, then it will use the default cache size and expiry. As the start of your `confluence/WEB-INF/classes/confluence-coherence-cache-config.xml` file you will notice the following:

```
1.<cache-mapping>
2.    <cache-name>*</cache-name>
3.    <scheme-name>default</scheme-name>
4.</cache-mapping>
```

So basically all caches will default to using the default scheme, which is defined as below:

```
1.<!-- Default scheme -->
2.<local-scheme>
3.    <scheme-name>default</scheme-name>
4.    <class-name>com.atlassian.confluence.cache.tangosol.ExpiryCountingLocalCache</class-name>
5.    <high-units>1000</high-units>
6.    <expiry-delay>3600</expiry-delay>
7.</local-scheme>
```

I.e. with a size of 1000 Objects and an expiry of 3600 seconds. Other schemes use the above as their default and either override the size of the cache, or the length of the expiry.

Common Schemes

In addition to the default scheme, there are also common schemes used in Confluence caches:

```
01.<!-- Common schemes -->
02.<local-scheme>
03.    <scheme-name>large</scheme-name>
04.    <scheme-ref>default</scheme-ref>
05.    <high-units>10000</high-units>
06.</local-scheme>
07.<local-scheme>
08.    <scheme-name>medium</scheme-name>
09.    <scheme-ref>default</scheme-ref>
10.    <high-units>5000</high-units>
11.</local-scheme>
12.<local-scheme>
13.    <scheme-name>small</scheme-name>
14.    <scheme-ref>default</scheme-ref>
15.    <high-units>100</high-units>
16.</local-scheme>
17.<local-scheme>
18.    <scheme-name>large-transient</scheme-name>
19.    <scheme-ref>default</scheme-ref>
20.    <high-units>10000</high-units>
21.    <expiry-delay>300s</expiry-delay>
22.</local-scheme>
23.<local-scheme>
24.    <scheme-name>user</scheme-name>
25.    <scheme-ref>default</scheme-ref>
26.    <high-units>5000</high-units>
27.    <expiry-delay>300s</expiry-delay>
28.</local-scheme>
```

RELATED TOPICS

Cache Performance Tuning
Confluence Cache Schemes
Cache Performance Tuning for Specific Problems
Requesting Performance Support
Administrators Guide
Configuration Guide

Configuring a Large Confluence Installation

Deploying *any* application to several thousand users requires care and planning, expecially if those users are going to be relying on the application to get their work done.

General Advice

Staged Rollout

Do not try to deploy Confluence immediately to your whole organisation. Instead, roll it out department by department, or project by project.

How Confluence will scale given a particular software and hardware configuration depends very much on how Confluence is likely to be used

in your organisation. Launching Confluence to everybody at once may seem like a neat idea, but it also means that any problems you might experience scaling the system up to your entire organisation will hit you *all at once*, annoy everyone and possibly hurt adoption.

Rolling Confluence out gradually will give you the chance to tune it as you go, resulting in a much more painless experience. There will also be organisational advantages: you can identify those teams or projects who are most likely to be successful 'early adopters', and those teams can experiment with how best a wiki might suit your organisation, and pass on their 'best wiki practices' as usage of Confluence expands.

Plugin Governance

Confluence plugins can add tremendous value. Before adding one, visit the plugin's page and explore its issues (available from the issue management link). Try the plugin in a test environment, and make sure to note any adverse effects after adding it to a production environment. Test plugins independently when upgrading.

Backup strategy

Disable the XML backup and use the Alternative Backup Strategy.

New Spaces Governance

For both performance and good practice, put some modest governance in place around the creation of new spaces, such as a simple request that includes a check for duplicates and some strategy around how to best use a space. Duplicates and unused spaces should be purged by a wiki gardener. Try to keep it to one space per group.

Choose User Management and Single Signon

It is possible to integrate with an LDAP repository or add a Single Signon solution later (especially with the addition of Crowd), but if possible it's best to configure this up front. You can configure access for only a specific group or set of groups, thereby keeping the gradual rollout.

Configuring your Application Server, Web Server and Database

Because Confluence can be deployed in so many server combinations, we do not currently have guides on the best tuning parameters for each individual server. We will be happy to provide support, however. If you have any tuning parameters that you find particularly useful for Confluence instances, feel free to share them with other Confluence users in the Confluence Community space.

Best Practices

Troubleshoot possible memory leaks

The Troubleshooting Confluence Hanging or Crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang. Many of the issues reported there are exacerbated with a large installation.

Memory Usage

The Java virtual machine is configured with a "maximum heap size" that limits the amount of memory it will consume. If Confluence fills up this maximum heap size it will run out of memory, and start behaving unpredictably. You can keep track of Confluence's memory usage from the System Information screen of the administration console:

Java VM Memory Statistics	
Total Memory	313 MB
Free Memory	140 MB
Used Memory	173 MB
Memory Graph	45 % Free
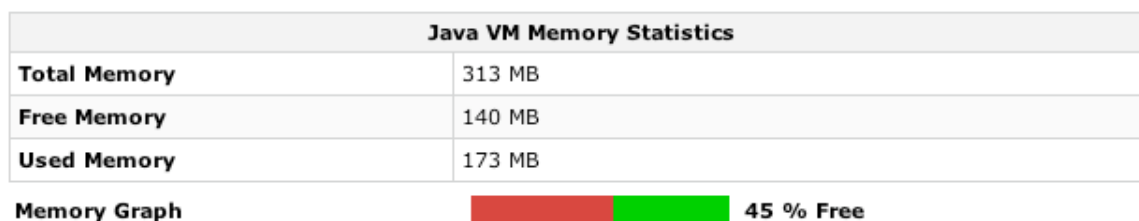
This example shows that, at the time of writing, confluence.atlassian.com is using 173MB of an allocated 313MB of heap. (The JVM was configured with a maximum heap size of 450MB, but this information is not available in the graph. The 313MB figure shows that the full 450MB of heap has not yet been needed)

Database Connection Pool

Confluence will need a database connection for each simultaneous user connection to the server. It is also a good idea to have 5-10 connections spare for Confluence internal processes such as backups, re-indexing or daily notification jobs.

Running out of pooled connections will cause the server to slow down as more users are waiting for a connection to be freed before starting their own request, and will eventually cause visible system errors as Confluence times out waiting for a database connection.

If you are using Confluence's internal connection pool, you can increase the number of available connections by modifying the `hibernate.c3p0.max_size` property in `{confluence_home}/confluence-cfg.xml`, and restarting Confluence. **Make sure** you have also configured your database to be able to support that many simultaneous connections.

Cache Sizes

The Performance Tuning page includes some useful rules of thumb for configuring the sizes of Confluence's internal caches.

RELATED TOPICS

Operating Large or Mission-Critical Confluence Installations
Performance Tuning
Confluence Clustering Overview
Requesting Performance Support
Administrators Guide
Configuration Guide

How Adaptavist Runs Confluence

Key information about performance tuning and how Adapatavist runs many instances of Confluence. This was posted to the mailing list but not captured in the forums so I wanted to copy it here:

Confluence Performance Recommendations from Adaptavist

Hi,
> 1) Opterons – for better or worse, we're solely an Intel shop at this point (i.e. I'll get some resistance going down the AMD path). It seems like the latest Woodcrest Xeons have caught up with the Opterons performance-wise (although AMD was smacking around Intel for a while before that). Have you the Woodcrest Xeons at all? (I'm guessing at this point you're probably just interested in pursuing an Opteron path but figured I'd ask.)
We found the Opterons are perfect for running multithreaded apps - especially when running lots of Confluence / JIRA installs on a server in their own Resin containers. In addition, the Opterons use a less electrical power than their Intel counterparts which reduces our energy bills - for that reason we've not tried the Woodcrest chip yet. The chassis used for Opteron chips is also more established than the Woodcrest counterpart and is thus cheaper at the moment and has more options. Any dual-core will likely give better results, but you'll obviously need your OS and other software on the server to be set-up to take advantage of them.
> 2) Resin vs. Tomcat – would you mind guessing at performance numbers from what you've seen? In a ton of Googling, it seems like people are saying that Tomcat was slow back with 4.x but got much faster with 5.x and even more with 5.5.x. See the comments in the first link.

Resin, properly configured, is still faster 🙂 In addition, it uses a lot less RAM (important when you have 50 web apps in their own containers). We spent a lot of time super-fine-tuning Resin (about 4 months if memory serves - huge thanks to Caucho and the chaps at BeJUG) to run Confluence very nicely indeed. I can't vouch for it's speed with regards to other apps. We've found Resin to be highly stable (when correctly configured) and it deals very well with that elusive Confluence memory leak (something you notice on a site like JavaPolis[1] with over 17,600 registered users). It's garbage collection, again when properly tuned, was better than Tomcat and we found many tasks easier to automate with Resin as compared to Tomcat.

Admittedly, a lot of the reasons that we chose Resin for are geared to an environment where we're running up to 50 Confluences on a single server, each in their own web app. That's quite a different scenario to what you are doing where you maybe have one or two Confluences on a server. Although, having said that, we use the same set-up for our dedicated hosting (we're hosting some real BIG Confluence installs as you'll know if you followed the discussion about the import routine we've been working on) and it works great in that environment too.

We use the commercial version of Resin[2] - it's much better than the OS version as it has fewer bugs, runs more smoothly and has some real nice features (read: absolutely critically essential for the sanity of our staff thus reducing our monthly bills for padded cells and therapy) for the type of environment we use it in. We also really liked Caucho's licensing of resin[3]: $500 per physical server with 2 cores (additional cores @ $500/core which is very reasonable) regardless of the number of Resin containers on that server.

It should be noted that some of the stats you provided links to were done on Windows running Cygwin - hardly an ideal server environment 😵 The second link (with all the graphs that people like me understand) was far more representative. We run on SuSE Linux 10 EMT64 (or something like that - whatever the latest version of their 64-bit OS is) so there's no Windows bloat getting in the way of the web apps, etc.
> 3) Memory. I think I'll go for DDR667 and see if I can bump Confluence up to 2 GB. Is there ever a point where you can allocate too much RAM? (i.e. java and/or Confluence just don't handle tracking that many cached objects well)
I think we're up to 16GB in most of our servers now. Confluence does enjoy more RAM (although containers such as Resin bring the overall RAM consumption down a fair bit - very noticeable on servers with 50 containers/apps). More RAM means more space to cache and longer gaps between forced GC. RAM allocation is also vital when it comes to the nightly backup (or "the dreadded backup hour" as we refer to it) - you can imagine the CPU and RAM spikes caused by 50 large Confluence installations all deciding to backup at the same time (roll on Confluence 2.3!)...

Should you have too much memory, you can always run a Quake server on there :o)

FYI: We also separate our database out on to a separate server.

Best Regards,

Guy

[1] http://www.javapolis.com - at last year's conference the Belgian's were somewhat annoyed at the term "SOA" which is an obscenity over there. They were also less than happy about the spoons in sexual positions plastered all over Antwerp (and several thousand Javapolians wearing the conference t-shirts). So this year Stephan and the crew have decided to push the boundaries to hitherto unimaginable levels - anyone who's seen the promo video will know exactly what I mean (and no, not the white painted bloke next to the urinal - the video goes waaaay beyond that - how they got James Gosling to... well, you'll have to wait and see)
[2] http://www.caucho.com/resin-3.0/features/overview.xtp
[3] http://www.caucho.com/sales/sales.xtp

Dan will probably be along in the morning to correct any mistakes I've made 🙂

-

Memory usage and requirements

Managing Confluence's performance and memory usage really depends on what resources are available - Confluence will run faster if you give it lots of memory for its caches, but it should still be able to run quite well in low-memory environments, with the right tuning. Below are some tips on getting the most out of your Confluence site:

Increasing the amount of memory available to Confluence

See Increasing JIRA Memory for details on how to increase the memory available to web application servers typically used to run Confluence.

Embedded Database

The embedded HSQL database that comes with Confluence essentially holds all your data in memory while the Confluence server is running. If you are running out of memory, you should consider migrating Confluence to some external RDBMS.

Caching

By default, Confluence keeps large in-memory caches of data to improve its responsiveness and the user experience. The trade off is an increase in memory requirements to support the cache. The default size of these caches will be reduced significantly in Confluence 1.5 (although this means that administrators of larger Confluence sites may need to configure the size of their caches to improve performance).

To customise confluences cache to meet your needs, see cache tuning.
To increase the amount of memory available to confluence, see Fix Out of Memory errors by Increasing Available Memory.

Mail error queue

Confluence keeps a copy of all emails that it failed to send within an internal error queue. In the event of intermittent failures such as network connectivity issues, the emails in this queue can be manually resent when the problem is fixed. Under certain circumstances, the mail queue can fill up with large objects. Since Confluence 1.4.1 the queue is regularly flushed, but if you get a *lot* of mail errors, you might get a spike in memory usage.

Attachments

The indexing of large attachments requires that the attachment be loaded into memory. In the case of large attachments, this can cause a temporary strain on the systems resources, and may result in indexing failing because the attachment could not be fully loaded into memory.

System backup / restore

The Confluence backup and restore process scales linearly with the size of data. This can have a significant impact on large Confluence instances where the amount of data exceeds the amount of available memory. If you are experiencing an `OutOfMemoryError` during either a backup or restore processes, then we strongly recommend that you choose and Alternative Backup Strategy.

If you do not wish to implement an alternative backup strategy, you could either switch to 'per space' backups or increase the memory allocated to Confluence.

If you encounter an `OutOfMemoryError` while restoring a backup and wish to overcome this issue by increasing memory, how much more will you need to make this process work? A good rule of thumb is to have a look at the size of the `entities.xml` file in your backup. This file contains all of the data Confluence will be loading, so at least that much is required. Add another 64-128Mb to ensure that Confluence has enough memory to load and function and that should be enough. This issue was resolved in Confluence post 1.4.x. To increase the amount of memory available to Confluence, see Fix Out of Memory errors by Increasing Available Memory.

Known issues that we do not have control over.

There are also some memory issues we don't have any control over. For example,

- There's a memory leak in the Oracle 10g JDBC drivers. Not much we can do about that.
- one customer found a rather nasty memory leak that appeared to originate inside Tomcat 5, but only using the IBM JDK on PowerPC.

If you are having problems that appear to result from a memory leak, file an issue on http://support.atlassian.com. Our memory profiler of choice is YourKit. It would be helpful to us if you can provide us with a memory dump from that tool showing the leak.

Confluence is taking long periods of time to respond to some actions

A common cause of random pauses in Confluence is the JVM running garbage collection. To determine if this is what is happening, enable verbose garbage collection and look at how long Java is taking to free up memory. If the random pauses match when Java is running its garbage collection, garbage collection is the cause of the pause.

Verbose garbage collection will generate log statements that indicate when Java is collecting garbage, how long it takes, and how much memory has been freed.

To enable gc logging, start Confluence with the option `-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:gc.log`. Replace `gc.log` with an absolute path to a `gc.log` file.

For example, with a Windows service, run:

```
tomcat5 //US//Confluence ++JvmOptions="-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:c:\confluence\logs\gc.log"
```

or in `bin/setenv.sh`, set:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:${CATALINA_BASE}/logs/gc.log"
```

If you modify `bin/setenv.sh`, you will need to restart Confluence for the changes to take effect.

What can you do to minimise the time taken to handle the garbage collection? See http://java.sun.com/docs/hotspot/gc1.4.2/ for details on tuning the JVM to minimise the impact that garbage collection has on the running application.

Requesting Performance Support

Basic Performance Troubleshooting Steps

Begin with the following procedures:

1. Go through the Troubleshooting Confluence Hanging or Crashing page to identify the major known performance problems
2. Proceed with the Performance Tuning tips to help optimize performance

Requesting Basic Performance Support

If those tips don't help or you're not sure where to start, open a support ticket starting with at least the basic information:

1. The atlassian-confluence.log
2. The catalina.out log (or your application server log), with a series of three thread dumps separated by 10 seconds
3. A description with as much detail as possible regarding:
   a. What changes have been made to the system?
   b. When did performance problems begin?
   c. When in the day do performance issues occur?
   d. What pages or operations experience performance issues?
   e. Is there a pattern?

Continue with as much of the Advanced Performance Troubleshooting information as you can.

Advanced Performance Troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

System Information

Confluence Server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Cache Statistics` (or save the page as HTML)
- Find out the exact hardware Confluence is running on
  - How many CPUs? What make and model? What MHz?
  - How much memory is installed on the machine?
  - How much memory is assigned to Confluence's JVM? (i.e. what are the -Xmx and -Xms settings for the JVM?)
  - What other applications are being hosted on the same box?

Confluence Content

- How many *users* are registered in Confluence?
- On average, to how many *groups* does each user belong?
- How many *spaces* (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many *pages*? (Connect to your database and perform `'select count(*) from content where prevver is null and contenttype = 'PAGE''`)
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform `'select count(*), sum(length(bandanavalue)) from bandana'`)

The Database

- What is the exact version number of Confluence's database server?
- What is the exact version number of the JDBC drivers being used to access it? (For some databases, the full filename of the driver JAR file will suffice)
- Is the database being hosted on the same server as Confluence?
- If it is on a different server, what is the network latency between Confluence and the database?
- What are the database connection details? How big is the connection pool? If you are using the standard configuration this information will be in your confluence_cfg.xml file. Collect this file. If you are using a Data source this information will be stored in your application server's configuration file, collect this data.

User Management

- Are you using external user management or authentication? (i.e. JIRA or LDAP user delegation, or single sign-on)
- If you are using external JIRA user management, what is the latency between Confluence and JIRA's database server?
- If you are using LDAP user management:
    - What version of which LDAP server are you using?
    - What is the latency between Confluence and the LDAP server?

Diagnostics

Observed Problems

- Which pages are slow to load?
    - If it is a specific wiki page, attach the wiki source-code for that page
- Are they always slow to load, or is the slowness intermittent?

Monitoring data

Before drilling down into individual problems, helps a lot to understand the nature of the performance problem. Do we deal with sudden spikes of load, or is it a slowly growing load, or maybe a load that follows a certain pattern (daily, weekly, maybe even monthly) that only on certain occasions exceeds critical thresholds? It helps a lot to have access to continuous monitoring data available to get a rough overview.

Here are sample graphs from the confluence.atlassian.com system, showing

Load
This graph shows the load for two consecutive days. The obvious pattern is that the machine is under decent load, which corresponds to the user activity, and there is no major problem.
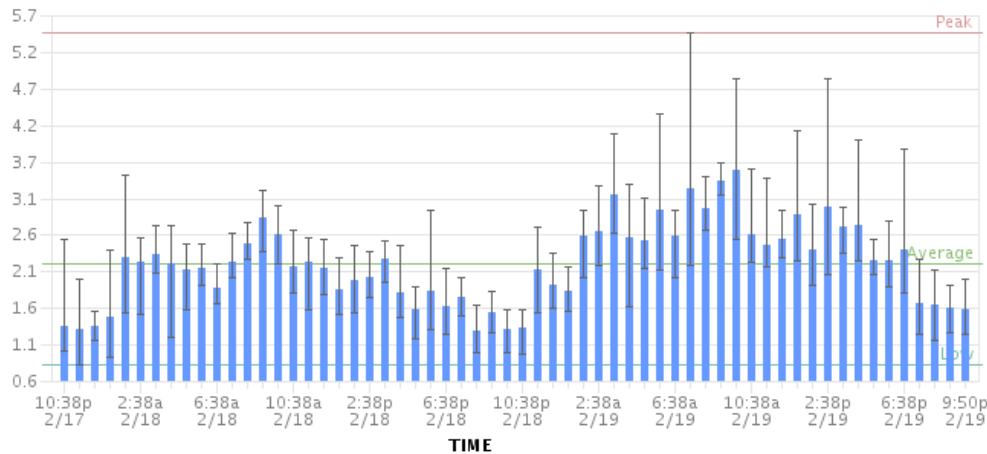
Resin Threads and Database Connections
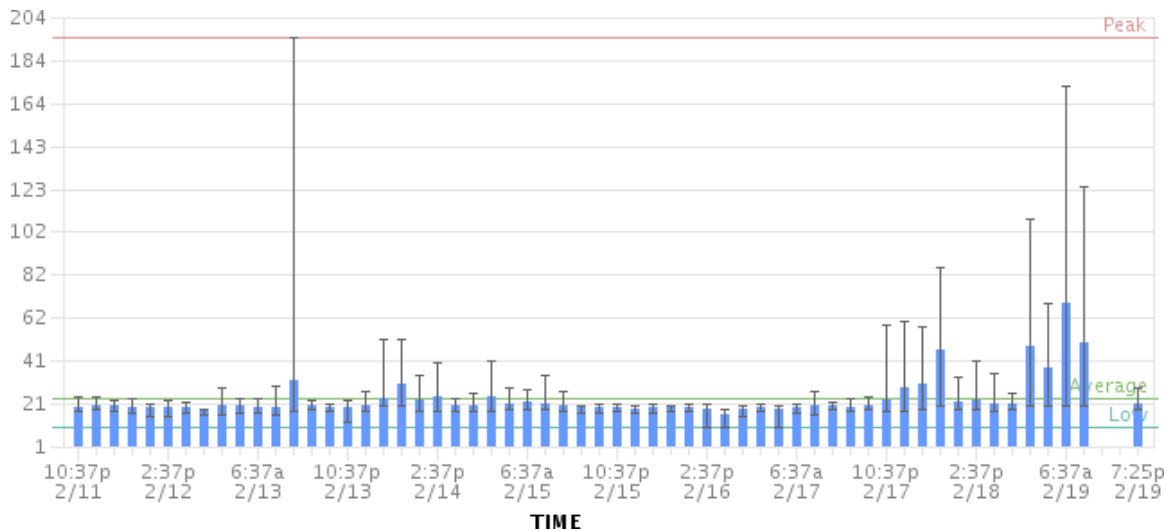
Active number of Java Threads

These two charts show the active threads in the application server (first chart) and the size database connection pool (second chart). As you can see, there was a sudden spike of server threads and a corresponding spike of db-connections.
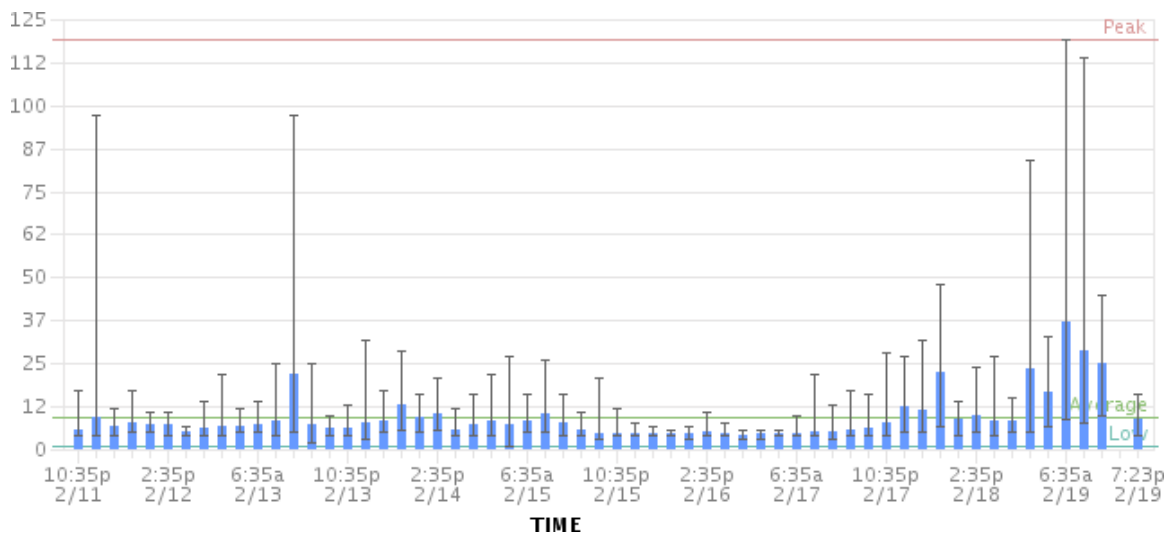
The database connection pool size

The database connection pool size peaked over 112, which happened to be more than the maximum number of connections the database was configured for (100). So it was no surprise that some requests to Confluence failed and many users thought it had crashed, since many requests could not obtain the crucial database connections.

We were able to identify this configuration problem quite easily just by looking at those charts. The next spikes were uncritical because more database connections were enabled.

The bottom line being: it helps a lot to monitor your Confluence systems continuously (we use Hyperic, for example), and it helps even more if you are able to send us graphs when you encounter problems.

Access logs

- How to audit Confluence - enabling user access logging, including redirecting the logs to a separate file
  - You can run this file through a log file analyser such as AWStats, or manually look through for pages which are slow to load.

Profiling and Logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using Troubleshooting Slow Performance Using Page Request Profiling.
  - If a single page is reliably slow, you should make several requests to that page
  - If the performance problem is intermittent, or is just a general slowness, leave profiling enabled for thirty minutes to an hour to get a good sample of profiling times
- Find Confluence's standard output logs (which will include the profiling data above). Take a zip of the entire logs directory.
- Take a thread dump during times of poor performance

CPU Load

- If you are experiencing high CPU load, please install the YourKit profile and attach two profiler dumps taken during a CPU spike. If the CPU spikes are long enough, please take the profiles 30-60 seconds apart. The most common cause for CPU spikes is a virtual

- machine operating system.
- If the CPU is spiking to 100%, try Live Monitoring Using the JMX Interface, in particular with the Top threads plugin.

Instance Metrics and Scripts

- It is essential to understand the user access and usage of your instance. Please use the access log scripts and sql scripts to generate Usage statistics for your instance.

Next Step

Open a ticket on https://support.atlassian.com and attach all the data you have collected. This should give us the information we need to track down the source of your performance problems and suggest a solution. Please follow the progress of your enquiry on the support ticket you have created.

If your site is non-responsive, please use our Live Support during business hours once you have created the ticket to escalate your problem.

Access Log Scripts

The access log scripts are attached to this page. To use the scripts:

1. Unzip the 7z file.
2. Copy all the daily access logs to a folder called `logs`.
3. Run `Atlassian-processDailyLog.rb`. This will generate a `csv` file called `summary.csv` and several directories which contain the access logs of each defined user action.
4. Run the appropriate script `Atlassian-processDailyLog-hourly.rb <admin/comment/create/edit/search/rss>`. Each script will generate a different `csv` file. For example, `Atlassian-processDailyLog-hourly.rb admin` will process the admin logs extracted in step 3.
5. Import the `csv` files to `www-log-Analysis.xls` (`summary.csv` to 'raw stats - daily' sheet and `admin.csv` to 'admin -hours' sheet, etc) to generate the load profiles and graphs. You may need to modify the number of rows in each sheet depending on the number of logs.

> **Note**
> All scripts are written in Ruby and assume the log file name contains the string 'confluence.atlassian.com-access.log'. Scripts need to be changed if another name is used. Modify the line: `filenameRegexp = Regexp.new('confluence.atlassian.com-access.log')`

Obtaining Confluence Instance Metrics

This page can be used as a guide to obtain detailed performance information of your instance.

Please read the Confluence Reporting HOWTO for information about the reporting capabilities of Confluence, including the `{sql}` macro, charting and security.

Users and usage

Users

What is the typical number of concurrent active users i.e. number of concurrent requests being processed?

- users with currently active requests
- users currently using Confluence: eg including reading a page, editing a page, viewing search results.
- users with sessions held in application server memory.
- users logged in *active users* (Note that Confluence uses "Remember Me" Session cookies and in my experience of Confluence, users never explicitly log out).
- define **user types** (viewer, editor, etc)

– light viewer
– rss reader
– searcher
– infrequent editor
– frequent editor
– administrator
– commenter

Usage

What is the average number of pages created per day, and similar usage stats (AWStat reports are a good starting place when User Access Logging is enabled)
To help interpret the raw access data, consider these important URL patterns:

Searches: `http://<host>/dosearchsite.action`
Rss requests: `http://<host>/createrssfeed.action`
Dashboard: `dashboard.action`
Creation: `createpage.action`
Editing: `http://<host>/pages/editpage.action`
Administrators: `http://<host>/admin/*`

Database usage statistics

⚠️ The following SQL statements are for PostgreSQL database. You may need to adjust the queries to suit your database.

Note: specify the date range

Table sizes

```
1. SELECT relname, reltuples, relpages FROM pg_class ORDER BY relpages DESC ;
```

- Example result:

relname	reltuples	relpages
pg_toast_1404472	115842	30116
trackbacklinks	451197	23832
bodycontent	170462	18197

The column `reltuples` is the number of rows in the table, `relpages` is the number of 8 KB pages used by the table. Indexes are included in this list as well.

In this example, the bodycontent table includes 170462 rows and is approximately 142 MB (18197 * 8 KB) in size.

Content created per day

```
1. select contenttype, min(number_of_changes), max(number_of_changes), avg(number_of_changes)
2. from (
3.        select contenttype, date_trunc('day', creationdate) , count(*) as number_of_changes
4.        from content
5.        where content.creationdate > date '2007-01-01' and version = 1
6.        group by contenttype, date_trunc('day', creationdate)
7.           ) as dates
8. group by contenttype
```

- Example result:

contenttype	min	max	avg
DRAFT	4	6	5.0000000000000000
MAIL	1	1	1.00000000000000000000
COMMENT	1	54	20.5040000000000000
USERINFO	1	45	15.8112449799196787
SPACEDESCRIPTION	1	3	1.1403508771929825
PAGE	1	119	21.4593495934959350
BLOGPOST	1	64	5.5925925925925926

Content edited per day

```
1. select contenttype, min(number_of_changes), max(number_of_changes), avg(number_of_changes)
2. from (
3.        select contenttype, date_trunc('day', lastmoddate) as changedate, count(*) as
number_of_changes
4.        from content
5.        where content.creationdate > date '2007-01-01'
6.        group by contenttype, date_trunc('day', lastmoddate)
7.           ) as dates
8. group by contenttype
```

- Example result:

contenttype	min	max	avg
BLOGPOST	1	718	14.4705882352941176
COMMENT	1	73	23.5120000000000000
DRAFT	4	6	5.0000000000000000
MAIL	1	1	1.00000000000000000000
PAGE	1	4658	130.2650602409638554
SPACEDESCRIPTION	1	4	1.2033898305084746
USERINFO	1	48	16.7991967871485944

Number of existing pages

```
1. select contenttype, count(*) from content group by content.contenttype
```

- Example result:

contenttype	count
MAIL	7914
COMMENT	12983
SPACEDESCRIPTION	232
DRAFT	10
PAGE	81465
USERINFO	13782
BLOGPOST	3308

Number of links per page

```
1. select http, max(linkcount), min(linkcount), avg(linkcount), stddev_pop(linkcount),
   stddev_samp(linkcount), var_pop(linkcount), var_samp(linkcount)
2. from
3. (
4.         select contentid, (links.destspacekey = 'http') as http, count(*) as linkcount
5.         from links group by contentid, (links.destspacekey = 'http')
6. ) as links_per_page
7. group by http
```

- Example result:

http	max	min	avg	stddev_pop	stddev_samp	var_pop	var_sa
false	1994	1	5.8366957470010905	32.7082672608353032	32.7104967872521825	1069.8307472062305489	1069.9
true	189	1	2.9633190883190883	6.3609167066017375	6.3614831031752836	40.4612613483250948	40.468

http	max	min	avg	stddev_pop	stddev_samp	var_pop	var_samp
f	600	1	5.8769371011850501	21.0221241817454213	21.0245200019346061	441.9297051127255987	442.03044131174
t	695	1	3.9395946999220577	12.8270107456162932	12.8282606370439241	164.5322046681558531	164.56427097193

Number of characters per content body

```
1. select max(blength), min(blength), avg(blength), stddev(blength), variance(blength)
2. from (select length(body) as blength from bodycontent) as bodylengths
3. where blength > 0
```

- Example result:

max	min	avg	stddev	variance
488707	1	2826.5649320388349515	8858.740996699238	78477292.046599816739

(Note this query takes a long time to execute.)

Number of characters per page body

```
1.select max(blength), min(blength), avg(blength), stddev(blength), variance(blength)
2.from (select length(bodycontent.body) as blength
3.            from bodycontent, content
4.            where bodycontent.contentid = content.contentid and contenttype='PAGE'
5.          ) as bodylengths
6.where blength > 0
```

- Example result:

max	min	avg	stddev	variance
488707	1	3333.0885906386048069	9884.337162920180	97700121.150284961908

Attachments

```
1.select count(*), max(filesize), min(filesize), avg(filesize), stddev(filesize), sum(filesize)
from attachments;
```

- Example result:

count	max	min	avg	stddev	sum
16082	107431588	0	207641.309725158562	2447904.9322	3339287543

Attachments per page

List the stats for attachments per page, only for those pages that actually have attachments.

```
1.select count(*) as pages_with_attachments, avg(attachments_per_page), max(attachments_per_page),
min(attachments_per_page), stddev(attachments_per_page) from
2.( select count(*) as attachments_per_page from attachments group by attachments.pageid ) as app
```

- Example result:

pages_with_attachments	avg	max	min	stddev
4197	3.8317846080533715	231	1	10.7013051235493489

Configuration / plugin data stored in Bandana

- Just the global context

```
1.select count(*), sum(length(bandanavalue)) from bandana where bandanacontext = '_GLOBAL'
```

- Example result:

count	sum
84	47729

- All of the information

```
1.select count(*), sum(length(bandanavalue)) from bandana
```

- Example result:

count	sum
665	153094

Content

It is essential to obtain the typical configuration of database (#pages, #spaces, #registered users, etc), based on Global Stats Plugin

Home directory usage statistics

On Unix-based environments like Linux and Mac OS X, you can use the following commands to gather information about the home directory usage.

Size of home directory components

```
du -sh /path/to/home/directory/*
```

- Example output:

```
5.9G attachments
4.0K backups
13M bundled-plugins
8.0K config
4.0K confluence.cfg.xml
216K fonts
3.4M framework-bundles
2.9G index
12M plugin-cache
114M plugins-temp
412K resources
4.0K restore
201M temp
84M thumbnails
222M viewfile
```

Number of attachments, including all versions

```
find /path/to/home/directory/attachments -type f | wc -l
```

Troubleshooting Slow Performance Using Page Request Profiling

This page tells you how to enable page-request profiling. With profiling turned on, you will see a record of the time it takes (in milliseconds) to complete each action made on any Confluence page. If Confluence is responding slowly, an internal timing trace of the slow page request can help to identify the cause of the delay.

You will need access to the Confluence server to view a profile.

In this page:

- Enabling Page-Request Profiling
- Profiling an Activity
- Example of a Profile
- Start Confluence with Profiling Enabled

Enabling Page-Request Profiling

From **Confluence 2.7**, you can use the '**Logging and Profiling**' option to enable or disable profiling.

You need to have System Administrator permissions in order to perform this function.

To enable page profiling,

1. Go to the '**Administration Console**' and click '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
2. The '**Logging and Profiling**' screen appears. Click the '**Enable Profiling**' button.
   If profiling is already enabled, the button will be labelled 'Disable Profiling' instead.

To disable page profiling,

1. Go to the '**Administration Console**' and click '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
2. The '**Logging and Profiling**' screen appears. Click the '**Disable Profiling**' button.
   If profiling is already disabled, the button will be labelled 'Enable Profiling' instead.

Screenshot: Changing Log Levels and Profiling

Performance Profiling

Profiling is currently OFF.

[Enable Profiling]

SQL Logging

[Enable SQL Logging]

Log4j Logging

Choose from one of the predefined logging options or configure logging below.

[Production]   [Diagnostic]

OR:

Customise specific logging settings

Add New Entry

Class/Package Name	New Level	
	INFO ▾	Add entry

Existing Levels

Class/Package Name	Current Level	New Level	
com.atlassian.confluence.cluster	INFO	INFO ▾	Remove
com.atlassian.confluence.cluster.safety	INFO	INFO ▾	Remove
com.atlassian.confluence.importexport.impl.PdfExporter	ERROR	ERROR ▾	Remove
com.atlassian.confluence.lifecycle	INFO	INFO ▾	Remove
com.atlassian.confluence.upgrade	INFO	INFO ▾	Remove
com.atlassian.core.util.FileUtils	ERROR	ERROR ▾	Remove
com.atlassian.upgrade	INFO	INFO ▾	Remove
net.sf.hibernate.cache.ReadWriteCache	ERROR	ERROR ▾	Remove
net.sf.hibernate.impl.SessionImpl	ERROR	ERROR ▾	Remove
net.sf.hibernate.type.CustomType	ERROR	ERROR ▾	Remove
net.sf.hibernate.util.JDBCExceptionReporter	ERROR	ERROR ▾	Remove
org.apache.fop	ERROR	ERROR ▾	Remove
root	WARN	WARN ▾	Remove

[Save]

Profiling an Activity

1. Enable profiling, using either of the methods described above.
   Profiles for every page hit, for all users, will now be logged to your application server's default logs until Confluence is restarted. Note that each time a user visits a link, a single profile is printed.
2. Confirm that profiles are being written to the Confluence log file — see Working with Confluence Logs for location of the log files and other details.
3. Perform the activity that is resulting in unusually slow response time.
4. Copy the profile for that action. When deciding which profiles to copy, look for the links that took a long time to respond. If a single page is slow, only that profile is necessary. If Confluence is generally or intermittently slow, copy all profiles logged during the slowdown until a reasonable sample has been collected.
5. If you were instructed to profile your instance by Atlassian technical support, attach all relevant profiles to your support ticket.
6. Turn profiling off again, using either of the methods described above.
7. Confirm that profiles are no longer being printed to the Confluence log file.

Example of a Profile

Below are the first few lines of a normal profile for accessing a page called Confluence Overview.

```
[344ms] - /display/ds/Confluence+Overview
  [313ms] - SiteMesh: parsePage: http://localhost:8080/display/ds/Confluence+Overview
    [313ms] - XW Interceptor: Before defaultStack: /pages/viewpage.action
(ViewPageAction.execute())
      [0ms] - SpaceAwareInterceptor.intercept()
      [16ms] - PageAwareInterceptor.intercept()
        [0ms] - AOP: PageManager.getPage()
        [16ms] - AOP: PermissionManager.hasPermission()
          [0ms] - AOP: SpacePermissionManager.hasPermission()
          [16ms] - AOP: SpacePermissionManager.hasPermission()
        [0ms] - AOP: SpacePermissionManager.hasPermission()
      [0ms] - AOP: SpacePermissionManager.hasPermission()
      [281ms] - XW Interceptor: After defaultStack: /pages/viewpage.action
(ViewPageAction.execute())
        [281ms] - XW Interceptor: After validatingStack: /pages/viewpage.action
(ViewPageAction.execute())
            ...
```

Start Confluence with Profiling Enabled

There may be some situations where you may wish to have Confluence profiling enabled during startup. This may be useful if you restart often and may forget to enable profiling for Support/Trouble-shooting purposes.

Edit the file *CONFLUENCE_HOME\confluence\WEB-INF\web.xml*. You should see a stanza similar to the one below. Set the parameter value for **autostart** to **true**:

```
<filter>
        <filter-name>profiling</filter-name>
        <filter-class>com.atlassian.core.filters.ProfilingAndErrorFilter</filter-class>
        <init-param>
            <!-- specify the which HTTP parameter to use to turn the filter on or off -->
            <!-- if not specified - defaults to "profile.filter" -->
            <param-name>activate.param</param-name>
            <param-value>profile</param-value>
        </init-param>
        <init-param>
            <!-- specify the whether to start the filter automatically -->
            <!-- if not specified - defaults to "true" -->
            <param-name>autostart</param-name>
            <param-value>true</param-value>
        </init-param>
    </filter>
```

Remember to turn it back to **false** or your logs will grow very large.

RELATED TOPICS

Requesting Performance Support
Working with Confluence Logs

Compressing an HTTP Response within Confluence

Confluence supports HTTP GZip transfer encoding. This means that if a user's web browser supports it, Confluence will compress the data it sends to the user. This will speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a

Confluence server.

ℹ Gzipping the HTTP Response is available in Confluence 1.4 and later.

You should turn on Confluence's GZip encoding if:

- Users are accessing Confluence over the Internet, or a WAN connection with limited bandwidth.
- You wish to reduce the amount of data transfer between the Confluence server and client.

If you are accessing Confluence over a Local Area Network or over a particularly fast WAN, you may wish to leave GZip encoding disabled. If the network is fast enough that transferring data from Confluence to the user isn't a limiting factor, the additional CPU load caused by having to compress each HTTP response may in fact slow Confluence down.

⛔ **Known issues in Confluence 2.7 and earlier**
There are known issues with the GZip filter and memory consumption evident in versions 2.7 of Confluence and earlier (CONF-9930). If you are running a large instance of Confluence 2.7 or earlier and frequently experiencing 'out of memory' errors, we recommend that you do not enable HTTP compression. These issues have been resolved in Confluence 2.8.

Enabling HTTP Compression

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Enable '**Compress HTTP Responses**'.

In Confluence 2.8 and later, you can configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement `urlrewrite-gzip-default.xml` file within the `WEB-INF/classes/com/atlassian/gzipfilter/` directory in your Confluence Installation Directory. A sample file is provided as an attachment. Generally speaking, it is unlikely that you will need to alter this file.

RELATED TOPICS

Performance Tuning
Administrators Guide

Performance Testing Scripts

Load Testing Confluence

This page contains scripts and hints on load-testing your Confluence installations.

Contents

Introduction

Before making a new Confluence instance available to your users it is useful to get a feel for how it will perform under your anticipated load and where you may need to consider improving your configuration to remove bottlenecks. Likewise, before making changes to your Confluence instance it would again be useful to assess the impact of these changes before making them live in a production context.

This kind of testing is not an exact science but the tools and process described here are intended to be a straightforward, configurable and extensible way of allowing you to begin this kind of load testing.

It will rarely be the case that these scripts will perform representative testing for you 'out of the box'. But either through configuration or by extending the scripts it should be possible to build an appropriate load test.

⛔ **Load testing scripts are not designed for a production environment**
The load testing scripts will update the data within the targeted Confluence instance and are not designed to be run against a production server. If you want to load test your production environment you will need to perform these tests on a backup of your data and restore your real data after the tests.

Setup

You will need the following -

- A Confluence server, set up and running with an admin user. The scripts assume a default username and password for this user: 'admin'/'admin'.
- Apache JMeter (currently version 2.3.4).
- The load testing scripts and resources which are available in our public Maven repository - version 3.1

> ℹ The Maven Repository has both ZIP and TAR archives. These archives contain the same files - if in doubt, download the ZIP file archive.

> ⛔ Users have reported problems when using the Windows built-in UNZIP utility. Please use a 3rd party unzip program such as WinZIP to extract these Performance Tests

The test scripts have been updated to work with Confluence 3.1 in version 3.1. Using an older version of the tests will result in errors when running the test.

Quick, Just Tell Me How To Run It.

If you don't want to read the rest of this document, here are the main points:

1. Create the test data:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts location>
-Jspace.zip=<path to a demo space ZIP file> \
-Jadmin.user=<username> -Jadmin.pass=<password>
```

2. Run the test:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-fixedload.jmx -Jscript.base=<scripts
location>
```

The remainder of this document is just an elaboration of those two steps.

> ⚠ For information on how to use JMeter please refer to the manual

Creating the Test Data

A known data set is required to run the testing against. By default this is the Confluence demo space (space key = DS) although this can be changed (more on this later). If you decide to use the Confluence demo space, ensure that the group "confluence-users" is able to update content in this space.

The script `jmeter-test-setup.jmx` is used to:

- create a set of users to be used in the test
- import the Confluence demo space for running tests against.

You should first ensure that you don't already have the demo space (key = DS) on your test instance. Delete it if you do.

Run the script from the `performance-testing` directory as follows:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts location>
-Jspace.zip=<path to a space export.zip> \
-Jadmin.user=<username> -Jadmin.pass=<password>
```

Where:

- `<scripts location>` is the absolute path to where you expanded the scripts e.g. `/Users/YourName/Download/performanceTest`. This defaults to the current directory. This is needed for the script to find its external resources and must be specified absolutely since JMeter occasionally does unexpected things with the working directory when it is running.

- `<path to a space export.zip>` is the absolute path to the space export zip you want to be used in your testing. For example, the path to `demo-site.zip` as found in your Confluence distribution or source: `<confluence`

```
install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup/demo-site.zip
```

- <username>> and <password> are the username and password for an admin user that is able to create Confluence users and to import spaces.

By default the setup process will create 250 users — 50 each of the following formats: tstreader<n>, tstcommentor<n>, tsteditor<n>, tstcreator<n> and tstsearcher<n>. The password for each matches the username.

A typical run of the setup script will only take a few seconds.

Removing the Test Data

You can reverse the effects of the setup script by setting the `remove.data` parameter to `true`, e.g.

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts base>
-Jremove.data=true -Jadmin.user=<username> -Jadmin.pass=<password>
```

Setup Script Parameters

You can modify the behaviour of the setup script via JMeter parameters. These are supplied on the command line in the form `-J<parameter name>=<parameter value>`.

Parameter	Default	Explanation
script.base	.	The absolute path to the script. Default to the current directory.
space.zip	N/A	The absolute path to space export zip file to be imported as test data.
remove.data	false	Run the script in reverse — remove all test data.
admin.user	admin	The admin user name used to import data and create users.
admin.pass	admin	The password for the admin user.
confluence.context	confluence	The confluence webapp context.
confluence.host	localhost	The address or host name of the test instance.
confluence.port	8080	The port of the test instance.
space.key	ds	The space key for the space import that will be tested against.
space.setup	true	Control whether the test space will be created (or removed).
commentor.max	50	The number of users to be created for making comments.
creator.max	50	The number of users to be created for adding pages.
editor.max	50	The number of users to be created for editing existing pages.
reader.max	50	The number of users to be created for viewing existing pages.
searcher.max	50	The number of users to be created for performing searches.

Setup Script Output

On the console you will see no obvious indication of success or otherwise. JMeter will output something similar to this:

```
Created the tree successfully
Starting the test @ Mon Apr 14 17:35:08 EST 2008 (1208158508222)
Tidying up ... @ Mon Apr 14 17:35:08 EST 2008 (1208158508928)
... end of run
```

The `scripts location/results` directory will contain the file `jmeter-result-setuptest.jtl`. There were failures or errors if there are any assertions in this file that have the value `true` for the failure or error element, e.g.

```
1.<assertionResult>
2.<name>Manage Users</name>
3.<failure>true</failure>
4.<error>false</error>
5.<failureMessage>Test failed: URL expected to contain /browseusers.action/</failureMessage>
6.</assertionResult>
```

Running the Test

The test script itself will put Confluence under a fixed load. Each thread group will attempt to do a certain amount of work for a prescribed period of time (30 minutes by default). This is by design so that load during test runs can accurately be compared against each other.

Execute the test as follows:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-fixedload.jmx -Jscript.base=<scripts location>
```

Where:
`<scripts location>` is the absolute path to where you extracted the scripts e.g. `/Users/YourName/Download/performanceTest`. This is needed for the script to find its external resources.

Test Behaviour

The test has a number of parameters to tweak its behaviour but generally speaking it has the rough format of:

- 5 groups of users - readers, commentors, searchers, editors and creators.
  - readers simply view a set of individual pages or browse space functionality.
  - commentors add comments to a set of pages.
  - searchers perform searches on a fixed set of keywords.
  - editors make small additions to the end of a set of pages.
  - creators add new pages to a particular space.
- Each individual user in each group will repeat for a fixed amount of time with a small pause between each request.

Note that there is **no execution of JavaScript** by the client. Keep this in mind if you use this test to gauge Confluence performance in a production environment.

There is also very little use of permissions in these tests. All data involved is accessible to all of the test users.

Test Script Parameters

You can modify the behaviour of the test script via JMeter parameters. These are supplied on the command line in the form `-J<parameter name>=<parameter value>`.

Parameter	Default	Explanation
script.base	.	The absolute path to the script. Defaults to the current working directory.
confluence.context	confluence	The confluence webapp context.
confluence.host	localhost	The address or host name of the test instance.
confluence.port	8080	The port of the test instance.
create.page.prefix	Nihilist	The title prefix for any created page e.g. Nihilist00001.
script.runtime	1800	The amount of time the script will run for in seconds.

Test Thread Parameters

Parameter	Default	Explanation
threads.reader	15	Number of readers.
pause.reader	2000	The approximate (within 500ms) millisecond pause between reader repeats.
threads.searcher	8	Number of searchers.
pause.searcher	2000	The approximate (within 500ms) millisecond pause between searcher repeats.
threads.creator	3	Number of page creators.
pause.creator	2000	The approximate (within 500ms) millisecond pause between creator repeats.
threads.editor	3	Number of page editors.
pause.editor	2000	The approximate (within 500ms) millisecond pause between editor repeats.
threads.commentor	4	Number of page commentors.
pause.commentor	2000	The approximate (within 500ms) millisecond pause between commentor repeats.

> ℹ️ In version 3.0 of the tests, it's now possible to control the percentage executions of certain actions. These percentages are defined in the "Thread Details" configuration screen.

So with the default parameters, you are emulating a load on Confluence of 33 concurrent users who will each be hitting the server

approximately every 2 seconds (16 users per second).

23 of these users are read only (searchers or readers) and 10 of them are read/write — 11 read only users per second and 5 read/write users per second.

Test Script Output

During the run of the test script Jmeter will output progress to the console of the form:

```
Created the tree successfully
Starting the test @ Fri Apr 18 00:07:39 EST 2008 (1208441259523)
Display Summary Results During Run + 462 in 77.6s = 5.9/s Avg: 1564 Min: 18 Max: 33738 Err: 1
(0.22%)
Display Summary Results During Run + 1338 in 189.9s = 7.0/s Avg: 3596 Min: 24 Max: 34545 Err: 0
(0.00%)
Display Summary Results During Run = 1800 in 257.6s = 7.0/s Avg: 3074 Min: 18 Max: 34545 Err: 1
(0.06%)
Display Summary Results During Run + 1046 in 200.9s = 5.2/s Avg: 4529 Min: 40 Max: 50461 Err: 0
(0.00%)
Display Summary Results During Run = 2846 in 438.2s = 6.5/s Avg: 3609 Min: 18 Max: 50461 Err: 1
(0.04%)
Display Summary Results During Run + 677 in 201.2s = 3.4/s Avg: 6638 Min: 46 Max: 27636 Err: 0
(0.00%)
Display Summary Results During Run = 3523 in 618.1s = 5.7/s Avg: 4191 Min: 18 Max: 50461 Err: 1
(0.03%)
Display Summary Results During Run + 561 in 197.5s = 2.8/s Avg: 8326 Min: 171 Max: 39494 Err: 0
(0.00%)
Display Summary Results During Run = 4084 in 798.3s = 5.1/s Avg: 4759 Min: 18 Max: 50461 Err: 1
(0.02%)
Display Summary Results During Run + 555 in 199.2s = 2.8/s Avg: 8247 Min: 160 Max: 45270 Err: 0
(0.00%)
Display Summary Results During Run = 4639 in 978.0s = 4.7/s Avg: 5177 Min: 18 Max: 504
```

Garbage Collector Performance Issues

> This document relates broadly to memory management with Sun's Hotspot JVM. These are recommendations based on Support's successful experiences with customers and their large Confluence instances.

Summary

- Set the Eden space up to 30-50% of the overall heap: -XX:NewSize=<up to half of your Xmx value>
- Use a parallel collector: `-XX:+UseParallelOldGC` (make sure this is Old GC)
- Disable explicit garbage collection triggered from distributed remote clients (`-XX:+DisableExplicitGC`)
- set the minimum and maximum Xmx and Xms values as the same (eg. `-Xms1024m -Xmx1024m`)
- Turn on GC logging (add the flags `-verbose:gc -Xloggc:<full-path-to-log> -XX:+PrintGCTimeStamps -XX:+PrintGCDetails`) and submit the logs in a support ticket
- Use Java 1.6
- Read below if heap > 2G

See Configuring System Properties for how to add these properties to your environment.

Background

Performance problems in Confluence generally manifest themselves in either:

- frequent or infrequent periods of viciously sluggish responsiveness, which requires a manual restart, or, the application eventually and almost inexplicably recovers
- some event or action triggering a non-recoverable memory debt, which in turn envelops into an application-fatal death spiral (Eg. overhead GC collection limit reached, or Out-Of-Memory).
- generally consistent poor overall performance across all Confluence actions

There are a wealth of simple tips and tricks that can be applied to Confluence, that can have a significantly tangible benefit to the long-term stability, performance and responsiveness of the application.

On this page:

- Summary
- Background
- Why this happens
- Appreciate how Confluence and the JAVA JVM use memory
- Memory is contiguous

- Figure out which (default) collector implementation your vendor is using
- Use the Parallel Garbage Collector
- Disable remote (distributed) garbage collection by JAVA clients
- Virtual Machines are Evil
- Use Java 1.6
- Use -server flag
- If using 64bit JRE for larger heaps, use CompressedOops
- Use NUMA if on SPARC, Opteron or recent Intel (Nehalem or Tukwila onwards)
- Use 32bit JRE if Heap < 2G
- JVM coredumps can be instigated by memory pressures
- Instigate useful Monitoring techniques
- Tuning the frequency of full collections
- Performance Tuning works

Why this happens

Confluence is basically a gel. Multiple applications, data-types, social networks and business requirements can be efficiently amalgamated together, leading to more effective collaboration. The real beauty of Confluence, however, is it's agility to mold itself into your organizations' DNA - your existing business and cultural processes, rather than the other way around - your organization having to adapt to how the software product works.

The flip side of this flexibility is having many competing demands placed on Confluence by it's users. Historically, this is an extraordinarily broad and deep set of functions, that really, practically can't be predicted for individual use cases.

The best mechanism to protect the installation is to place Confluence on a foundation where it is fundamentally more *resilient* and able to react and cope with competing user requirements.

Appreciate how Confluence and the JAVA JVM use memory

The java memory model is naive. Compared to a unix process, which has four *intensive* decades of development built into time-slicing, inter-process communication and intelligent deadlock avoidance, the JAVA thread model really only has 10 years at best under it's belt. As it is also an interpreted language, particular idiosyncrasies of the chosen platform Confluence is running can also influence how the JRE reacts. As a result it is sometimes necessary to *tune* the jvm parameters to give it a "hint" about how it should behave.

There are circumstances whereby the JAVA jvm will take the 'mediocre' option in respect to resource contention and allocation, and, 'struggle' along with ofttimes highly impractical goals. For example, The JRE will be quite happy to perform at 5 or 10% of optimum capacity if it means overall application stability and integrity can be ensured. This often translates into periods of extreme sluggishness, which effectively means that the application isn't stable, and isn't integral (as it cannot be accessed).

This is mainly because JAVA shouldn't make assumptions on what kind of runtime behavior an application needs, but, it's plain to see that the charter is to assume 'business-as-usual' for a wide range of scenarios and really only react in the case of 'dire' circumstances.

Memory is contiguous

The Java memory model *requires* that memory be allocated in a *contiguous* block. This is because the heap has a number of side data structures which are indexed by a scaled offset (ie n*512 bytes) from the start of the heap. For example, updates to references on objects within the heap are tracked in these "side" data structures.

Consider the differences between:

1. Xms (the *allocated* portion of memory)
2. Xmx (the *reserved* portion of memory)

Allocated memory is fully backed, memory mapped physical *allocation* to the application. That application now owns that segment of memory.

Reserved memory (the difference between Xms and Xmx) is memory which is *reserved* for use, but not physically mapped (or backed) by memory. This means that, for example, in the 4G address space of a 32bit system, the *reserved* memory segment can be used by other applications, but, because JAVA requires *contiguous* memory, if the *reserved* memory requested is occupied, the OS must swap that memory out of the reserved space either to another non-used segment, or, more painfully, it must swap to disk.

Permanent Generation memory is also contiguous. The net effect is even if the system has vast quantities of *cumulative* free memory, Confluence demands *contiguous* blocks, and consequently undesirable swapping may occur if segments of requested size do not exist. See Causes of OutOfMemoryErrors for more details.

Please be sure to position Confluence within a server environment that can successfully complete competing requirements (operating system, contiguous memory, other applications, swap, and Confluence itself).

Figure out which (default) collector implementation your vendor is using

Default JVM Vendor implementations are subtly different, but in production can differ enormously.

SUN *by default* splits the heap into three *spaces*

1. Eden (Nursery, or Scavenger)
2. Tenured (Old)
3. Permanent Generation (classes & library dependencies)

Objects are central to the operation of Confluence. When a request is received, the Java runtime will create new objects to fulfill the request in the Eden Space. If, after some time, those objects are still required, they may be moved to the Tenured (Old) space. But, typically, the *overwhelming majority* of objects created die young, within the Eden Space. These are objects like method local references within a while or for loop, or Iterators for scanning through Collections or Sets.

But in IBM J9 the default policy is for a single, contiguous space - one large heap. The net effect is that for large Websphere environments, garbage collection can be terribly inefficient - and culpable to suffer outages during peak periods.

> ⚠️ For *larger instances with performance issues*, it is recommended to tune Confluence such that there is a large Eden space, at *up to* 50% the overall size of the heap.

`-XX:NewSize=XXXm` where XXX is the size in megabytes, is the command line parameter. `-XmnXXXm` can also be used interchangeably. Ie. -XX:NewSize=700m, -Xmn700m

By setting a larger NewSize, the net effect is that the JRE will spend less time garbage collecting, clearing dead memory references, compacting and copying memory between spaces, and more time *doing actual work*.

Use the Parallel Garbage Collector

Confluence out of the box, and Sun Java as default, uses the *serial* garbage collector on the Full Tenured heap. The Eden space is collected in parallel, but the Tenured is not. This means that at a time of load, if a full collection event occurs, since the event is a 'stop-the-world' serial event, *all application threads* other than the garbage collector thread are taken off the CPU. This can have severe consequences if requests continue to accrue during these 'outage' periods. As a rough guide, for every gigabyte of memory allocated, allow a full second (exclusive) to collect.

If we parallelize the collector on a multi-core/multi-cpu architecture instance, we not only *reduce* the total time of collection (down from whole seconds to fractions of a second) but we also improve the *resiliency* of the JRE in being able to recover from high-demand occasions.

Additionally, Sun provide a CMS, Concurrent Mark-Sweep Collector (-XX:+UseConcMarkSweepGC), which is more optimized for higher-throughput, server-grade instances. As a general rule the Parallel Collector (-XX:+UseParallelOldGC) is suitable for most installations, if you are keen to extract the best performance available, then the CMS collector is an option.

Disable remote (distributed) garbage collection by JAVA clients

Many clients integrate 3rd party, or their own custom applications to interrogate, or add content to Confluence via it's RPC interface. The Distributed Remote Garbage Collector in the client uses RMI to trigger a remote GC event in the server, Confluence. Unfortunately, as of this writing, a System.gc() call via this mechanism triggers a full, serial collection of the entire Confluence heap (as it needs to remove references to remote client objects in it's own deterministic object graph). This is a deficiency in the configuration and/or implementation of the JVM, and, has the potential to impact severely if the remote client is poorly written, or operating within a constricted JVM.

This can be disabled by using the flag `-XX:+DisableExplicitGC` at startup.

Virtual Machines are Evil

Vmware Virtual Machines, whilst being extremely convenient & fantastic, also obviously being a very strong growth segment of the industry, cause particular problems for JAVA applications because it's very easy for host operating system resource constraints (ie, temporarily insolvent memory availability, or I/O swapping), to cascade into the JAVA VM and manifest itself in extremely unusual, frustrating and seemingly illogical problems. We already document some disk I/O metrics with VMware images, and, although we now *officially* support the use of virtual instances, we absolutely do not recommend them unless maintained correctly.

This is not to say that vmware instances cannot be used, but, they must be used with due care, proper maintenance and configuration. Besides, if you are reading this document because of poor performance, the first port of call should be to remove any virtualization - emulation will never beat the real thing, and at the least introduces black box variability into the system.

Use Java 1.6

Java 1.6 is generally regarded via public discussion to have an approximate 20% performance improvement over 1.5. Indeed, our own internal testing revealed this statistic to be credible. 1.6 is compatible for all supported versions of Confluence, and we **strongly recommend** that installations not using 1.6 migrate.

Use -server flag

The hotspot server JVM has specific code-path optimizations which yield an approximate 10% gain over the client version. Most installations *should* already have this selected by default, but it is still wise to force it with -server, especially on some Windows machines.

If using 64bit JRE for larger heaps, use CompressedOops

For every JDK release Sun also build a "Performance" branch in which specifically optimized performance features can be enabled; It is available on the Sun Java SE page after a brief survey. These builds are certified production grade.

Some blogs have suggested a 25% performance gain AND reduction in Heap size when using this parameter. The use and function of the `-XX:+UseCompressedOops` parameter is more deeply discussed on Sun's Official Wiki (Co-incidentally, using Confluence!)

Use NUMA if on SPARC, Opteron or recent Intel (Nehalem or Tukwila onwards)

-XX:+UseNUMA flag enables the java heap to take advantage of Non-Uniform-Memory-Architectures. JAVA will place data structures relevant to the thread which it owns / operates on, in memory locations closest to that particular processor. Depending on the environment, gains can be substantial. Intel market NUMA as Quick Path Interconnect™.

Use 32bit JRE if Heap < 2G

Using a 64bit JRE when the heap is under 2G will hurt substantially in heap and performance, as nearly every object, reference, primitive, class and variable will cost twice as much to address.

A 64bit JRE/JDK is only recommended if heaps greater than 2G are required. If so, use CompressedOops.

JVM coredumps can be instigated by memory pressures

If your instance of Confluence is throwing Java coredumps, it's known that memory pressure and space / generation sizings can influence the frequency and/or occurrence of this phenomena.

If your Tomcat process completely disappears and the logs record similar to:

```
#
# An unexpected error has been detected by HotSpot Virtual Machine:
#
#  SIGSEGV (0xb) at pc=0xfe9bb960, pid=20929, tid=17
#
# Java VM: Java HotSpot(TM) Server VM (1.5.0_01-b08 mixed mode)
# Problematic frame:
# V  [libjvm.so+0x1bb960]
#

--------------  T H R E A D  ---------------

Current thread (0x01a770e0):  JavaThread "JiraQuartzScheduler_Worker-1" [_thread_in_vm, id=17]

siginfo:si_signo=11, si_errno=0, si_code=1, si_addr=0x00000000

Registers:
 O0=0xf5999882 O1=0xf5999882 O2=0x00000000 O3=0x00000000
 O4=0x00000000 O5=0x00000001 O6=0xc24ff0b0 O7=0x00008000
 G1=0xfe9bb80c G2=0xf5999a48 G3=0x0a67677d G4=0xf5999882
 G5=0xc24ff380 G6=0x00000000 G7=0xfdbc3800 Y=0x00000000
 PC=0xfe9bb960 nPC=0xfe9bb964
```

We recommend a bug report be submitted to the JAVA Vendor, but in Support's experience, **resource contention issues and core dumps are tightly coupled**.

Instigate useful Monitoring techniques

At all times, the best performance tuning recommendations are based on current, detailed metrics. This data is easily available and configurable, and helps us at Atlassian, **tremendously**, when diagnosing reported performance regressions.

1. enable JMX monitoring
2. enable Confluence Access logging
3. enable Garbage Collection Logging
4. Take Thread dumps at the time of regression. If you can't get into Confluence, you can take one externally.
5. Jmap can take a memory dump in real time without impacting the application. Syntax: `jmap -heap:format=b <process_id>`

Great tools available include:

- The very excellent VisualVM, Documentation.
- Thread Dump Analyzer - great all round thread debugging tool, great for identifying deadlocks.
- Samurai, an excellent alternative thread analysis tool, great for iterative dumps over a period of time.
- GC Viewer - getting a bit long in the tooth, but is a good mainstay for GC analysis.
- GChisto - A new GC analysis tool written by members of the Sun Garbage Collection team.

Documentation:

- Sun's White Paper on Garbage Collection in JAVA 6.
- Sun's state-of-the-art JavaOne 2009 session on garbage collection (registration required).
- IBM stack: Java 5 GC basics for WebSphere Application Server.
- An Excellent IBM document covering native memory, thread stacks, and how these influence memory constricted systems. Highly recommended for additional reading.
- The complete list of JRE 6 options

> ⚠️ Atlassian recommends at the very least to get VisualVM up and running (You need JMX), and to add Access and Garbage Collection logging.

Tuning the frequency of full collections

The JVM will generally only collect on the full heap when it has no other alternative, because of the relative size of the Tenured heap (it is typically larger than Eden), and the natural probability of objects within tenured not being eligible for collection. (ie. they are still alive).

Some installations can trundle along, only ever collecting in Eden space. As time goes on, some object will survive the initial Eden object slaughterhouse and be promoted to Tenured. At some point, it will be dereferenced and no longer reachable by the deterministic, directed object graph. However, the occupied memory will still be held in limbo as "dead" memory until a collection occurs in the Tenured space to clear and compact the space.

It is not uncommon for moderately sized Confluence installations to reclaim as much as 50% of the current heap size on a full collection; This is because full collections occur so infrequently. By reducing the occupancy fraction heap trigger, this means that more memory will be available at any time, meaning less swapping / object collections will occur *when needed*, ie, during the busy hour.

Atlassian would classify frequency tuning on collections as an **advanced** topic for further experimentation, and is provided for informational purposes only. Unfortunately, it's impractical for Atlassian to support this degree of 'ergonomics'.

Performance Tuning works

Atlassian has a number of high profile and some *extremely* high demanding, mission-critical clients who have successfully, through some trial and error, applied these recommendations to production instances and have very positively improved their Instances. For any more information, or guidance, please lodge a support case with the relevant information.

Scheduled Jobs

This page provides a quick overview of the jobs that are scheduled to run regularly in your Confluence instance.

Job Name	Description
backupJob	performs a site backup
mailQueueFlushJob	sends notifications that have been queued up
referralQueueFlushJob	referrals to Confluence pages are queued up. This job writes this referrals to the database
taskQueueFlushJob	flushes the task queue
cleanTempDirectoryJob	this cleans up temp files created in Confluence home temp directory (created by exports etc.)
dailyReportJob	sends out an email summary of all changes in Confluence to all subscribers
clearOldMailErrorsJob	notifications that fail to send due to errors are added to the mail error queue. This job resets this clear periodically.
indexQueueFlushJob	each content update to Confluence needs to be updated in index so search results are accurate. This job flushes changes to the index.
indexOptimizerJob	index optimization is performed to compact the index and maintain searching performance. This task is expensive and does not need to be performed too regularly. If you see Confluence performance deteriorate around 3pm, you can try scheduling this job for 3am only and check that search performance remains reasonable.
indexQueueCleanJob	this job is responsible for periodically triggering an Index Queue clean to ensure that size of the index queue does NOT grow indefinitely.
mailPollJob	polls POP accounts on all spaces that have them configured.
clusterSafetyJob	ensures that only one cluster is ever writing to the database at one time. For non-clustered instances, this job is still useful for alerting customers that have accidentally deployed two instances of Confluence against the same database.

Search

📄 Setup External Search Tool To Index Confluence

📄 Setup Confluence To Index External Sites

Setup Confluence To Index External Sites

Confluence Indexing External Sites

Confluence cannot easily index external sites due to technical reasons, but there are two alternatives:

1. Embed External Pages Into Confluence
2. Replace Confluence Search

Technical Reasons

Confluence indexes pages using a customised Lucene search engine that returns matching pages, mail and blog posts for which the searcher has view permission. It would require significant source code modifications to enable Confluence to process search results from external pages, as the indexing process has been customised to utilise internal Confluence metadata. Note that users can still index content from new attachment filetypes.

Embed External Pages Into Confluence

If you only have a small number of external sites to index, you may prefer to enable the HTML-include Macro and use it embed the external content inside normal Confluence pages.

Replace Confluence Search

Use your own programmer resources to replace Confluence's internal search with a crawler that indexes both Confluence and external sites. This advanced option is easier than modifying the internal search engine. It requires removing Confluence internal search from all pages and replacing the internal results page with your own crawler front-end.

1. Setup a replacement federated search engine to index the Confluence site, as well as your other sites, and provide the results that way. You would need to host a web crawler, such as these open-source crawlers. Note that you can perform a search in Confluence via the remote API

2. Replace references to the internal search by modifying the site layout so that it links to your search front-end

3. Host another site containing the search front-end. You may wish to insert it into a suitable context path in your application server so that it appears to be from a path under Confluence. Tomcat sets Confluence's paths from the Confluence install\confluence\WEBINF\web.xml file.

RELATED TOPICS

Setup External Search Tool To Index Confluence

Setup External Search Tool To Index Confluence

Any web crawler can be configured to index Confluence content, for example the Google Search Appliance or similar. If a login is required to view content that will be indexed, you should create a Confluence user specifically for the search crawler to use. Grant this user view rights to all content you wish to index, but deny that user all delete and administration rights. This ensures that an aggressive crawler will not be able to perform actions that could modify the site. There is also a forum thread on Google Mini integration.

External applications can also use the search function in the Confluence Remote API.

Related Information

📄 Setup External Search Tool To Index Confluence

📄 Setup External Search Tool To Index Confluence

📄 Setup Confluence To Index External Sites

📄 Setup Confluence To Index External Sites

📄 Setup External Search Tool To Index Confluence

📄 Setup Confluence To Index External Sites

📄 Setup External Search Tool To Index Confluence

📄 Setup Confluence To Index External Sites

📄 Setup Confluence To Index External Sites

   📄 Setup External Search Tool To Index Confluence

   📄 Setup External Search Tool To Index Confluence

   📄 Setup Confluence To Index External Sites

   📄 Setup External Search Tool To Index Confluence

   📄 Setup Confluence To Index External Sites

   📄 Integrate Confluence Search to Jira Search

   📄 Indexing External Content

Security

- Adding SSL for Secure Logins and Page Security
- Anonymous Access to Remote API
- Confluence Cookies
- Enabling or Disabling Public Signup
- Hiding External Links From Search Engines
- Hiding the People Directory
- Managing External Referrers
    - Excluding external referrers
    - Hiding external referrers
    - Ignoring External Referrers
- User Email Visibility
- Spam Prevention via Captcha

Adding SSL for Secure Logins and Page Security

This document describes how to configure Confluence to use a HTTPS encrypted secure socket layer for user logins and page data.

Unencrypted confidential data within Confluence may be intercepted by an attacker. To secure user logins, you can enable access via HTTPS (HTTP over SSL), and require its use for pages where passwords are sent. In some cases where issue data is sensitive, all pages can be set to be accessed over HTTPS.

Enabling SSL access is different for each application server, but specifying which pages to require protection for is generic. This document is specific to Tomcat, the default application server shipped with Confluence.

On this page:

- Adding Secure User Logins
    - Creating A New SSL Certificate
    - Verify the Certificate is in the Correct Location
    - Specifying URL Patterns to be Redirected
- Troubleshooting

Adding Secure User Logins

Adding HTTPS requires a valid SSL certificate. If you have a Certificate prepared, skip to the 'Modify the `<INSTALL>/conf/server.xml` File' section.

Creating A New SSL Certificate

Creating a self-signed certificate

> ℹ️ The following commands are in reference to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to this document.

On Windows, perform the following at the command prompt:

```
1."%JAVA_HOME%\bin\keytool" -genkey -alias tomcat -keyalg RSA
```

Or on other platforms, perform the following at the command prompt:

```
1.$JAVA_HOME/bin/keytool -genkey -alias tomcat -keyalg RSA
```

Some questions will be asked, including a password for the certificate (the default is 'changeit'). Please note down what you choose, as it will be used in the next step.

> 🚫 **"IE7 on Vista Issue"**
> If your clients will access Confluence from Internet Explorer 7 on Vista, please ensure that you specify the -keyalg RSA flag. By default the SHA1 algorithm is used, which results in 'Internet Explorer cannot display the webpage' errors on IE7 on Vista. Apparently on JDK 1.6 you also need to specify the -sigalg MD5withRSA flag since -keyalg RSA will still result in SHA1 being used (see this blogpost for more information).

Modify the `<INSTALL>/conf/server.xml` File

In the confluence directory, open the *conf/server.xml* file and insert one of the following just after the closing `</Engine>` tag:

1. For users of Confluence 2.10 or later:
   Open `conf/server.xml`, uncomment the lines:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                        maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
3.                        enableLookups="false" disableUploadTimeout="true"
4.                        acceptCount="100" scheme="https" secure="true"
5.                        clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
6.                        URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"/>
```

   Or for users of Confluence 2.2 to 2.9.2:
   Open `conf/server.xml`, uncomment the lines:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                                        maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
3.                                        enableLookups="false" disableUploadTimeout=
"true"
4.                                        acceptCount="100" scheme="https" secure=
"true"
5.                                        clientAuth="false" sslProtocol="TLS"
6.                                        URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>" />
```

> ⚠️ If your Confluence server is running off Apache Tomcat version 6.0.0 or later, you should ensure that the parameter-value pair `SSLEnabled="true"` has been added to the `Connector` tag above.

Establishing a CA-issued Certificate

In preparation for a production instance, an official CA-issued key pair is required. Find instructions in the Tomcat documentation.

Verify the Certificate is in the Correct Location

By default, Tomcat will look for the certificates in the file `C:\Documents and Settings\\#CURRENT_USER#\.keystore` on Windows or `~/.keystore` on Unix. If your Certificate is not in this location, you will need to update your `<INSTALL>/conf/server.xml` file as outlined below, so that Tomcat can find it.

1. For users of Confluence 2.2 or later:
   Open `conf/server.xml`, add the `keystoreFile="<MY_CERTIFICATE_LOCATION>"` parameter to the Connector tag as shown below:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                                        maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
3.                                        enableLookups="false" disableUploadTimeout=
"true"
4.                                        acceptCount="100" scheme="https" secure=
"true"
5.                                        clientAuth="false" sslProtocol="TLS"
6.                                        URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>" keystoreFile="<MY_CERTIFICATE_LOCATION>" />
```

> ⚠️ Make sure to change your Server Base URL to https.

Specifying URL Patterns to be Redirected

Restart Tomcat and access your instance on `https://<MY_BASE_URL>:8443/`.

For more detailed information on setting up SSL with Tomcat (including additional configuration options), have a look at Tomcat 5.5 SSL or Tomcat 6 SSL.

Although HTTPS is now activated and available, the old HTTP URLs (http://localhost:8080) are still available. In most situations one wants these URLs to continue working, but for some to redirect to their HTTPS equivalent.

> ⚠️ If you have changed the port that the SSL connector is running on from the preconfigured value of 8443, you must update the `redirectPort` attribute of the standard HTTP connector to reflect the new SSL port. Tomcat needs this information to know which port to redirect to when an incoming request needs to be secure.

If security is a concern, we recommend using SSL encryption site wide, for the reasons listed here: CONF-4116. To do this:

Edit the `confluence/WEB-INF/web.xml` file and add the following declaration to the end, **before** the `</web-app>` tag:

```
01.<security-constraint>
02.    <web-resource-collection>
03.        <web-resource-name>Restricted URLs</web-resource-name>
04.        <url-pattern>/</url-pattern>
05.    </web-resource-collection>
06.    <user-data-constraint>
07.        <transport-guarantee>CONFIDENTIAL</transport-guarantee>
08.    </user-data-constraint>
09.</security-constraint>
```

Once this change is made, restart Confluence and access http://localhost:8080. You should be redirected to https://localhost:8443/login.action.

If you'd like to protect login.action only:

> ℹ️ Please note that redirecting login.action only will leave the rest of the site open to http connections, but will not redirect the user back to http once they're on https. To do that, you'll need to add a rewrite rule in Apache or IIS.

```
01.<security-constraint>
02.    <web-resource-collection>
03.        <web-resource-name>Login and Restricted Space URLs</web-resource-name>
04.        <url-pattern>/login.action</url-pattern>
05.    </web-resource-collection>
06.    <user-data-constraint>
07.        <transport-guarantee>CONFIDENTIAL</transport-guarantee>
08.    </user-data-constraint>
09.</security-constraint>
```

Note that the example above specifies a `url-pattern` for the login URL */login.action*. This means that whenever a user tries to access the unprotected version of the login page, they will be redirected automatically to the secured version of it.

If you want to protect individual spaces, there isn't a complete way of doing this at the moment. You can add a pattern like this:

```
01.<security-constraint>
02.    <web-resource-collection>
03.        <web-resource-name>Login and Restricted Space URLs</web-resource-name>
04.        <url-pattern>/login.action</url-pattern>
05.        <url-pattern>/display/SALARIES/*</url-pattern>
06.    </web-resource-collection>
07.    <user-data-constraint>
08.        <transport-guarantee>CONFIDENTIAL</transport-guarantee>
09.    </user-data-constraint>
10.</security-constraint>
```

Troubleshooting

Check the Confluence Knowledge Base articles at Troubleshooting SSL.

Anonymous Access to Remote API

Sites may wish to disable anonymous access to the remote API to make it harder for malicious users to write 'bots' that perform bulk changes to the site. If you wish to enable the Remote APIs but do not want anonymous users to access Confluence remotely, you can disable anonymous access from the **Administration Console**.

To disable anonymous access to Remote APIs,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
4. Select '**Off**' beside '**Anonymous Access to API**'.
5. '**Save**' your changes.

RELATED TOPICS

Enabling or Disabling Public Signup

Adding SSL for Secure Logins and Page Security

Hiding external referrers

Excluding external referrers

Hiding External Links From Search Engines

Hiding the People Directory

User Email Visibility

Anonymous Access to Remote API

Ignoring External Referrers

Spam Prevention via Captcha

Managing External Referrers

Administrators Guide Home Confluence Documentation Home

Confluence Cookies

Confluence uses Seraph, an open source framework, for HTTP cookie authentication.

Cookies

Confluence uses two cookies. The first, a JSESSIONID cookie, is created by the Application Server and used for session tracking purposes. The second, the 'Remember my login on this computer' cookie, is generated by Confluence when the user selects the check-box on the log-in page.
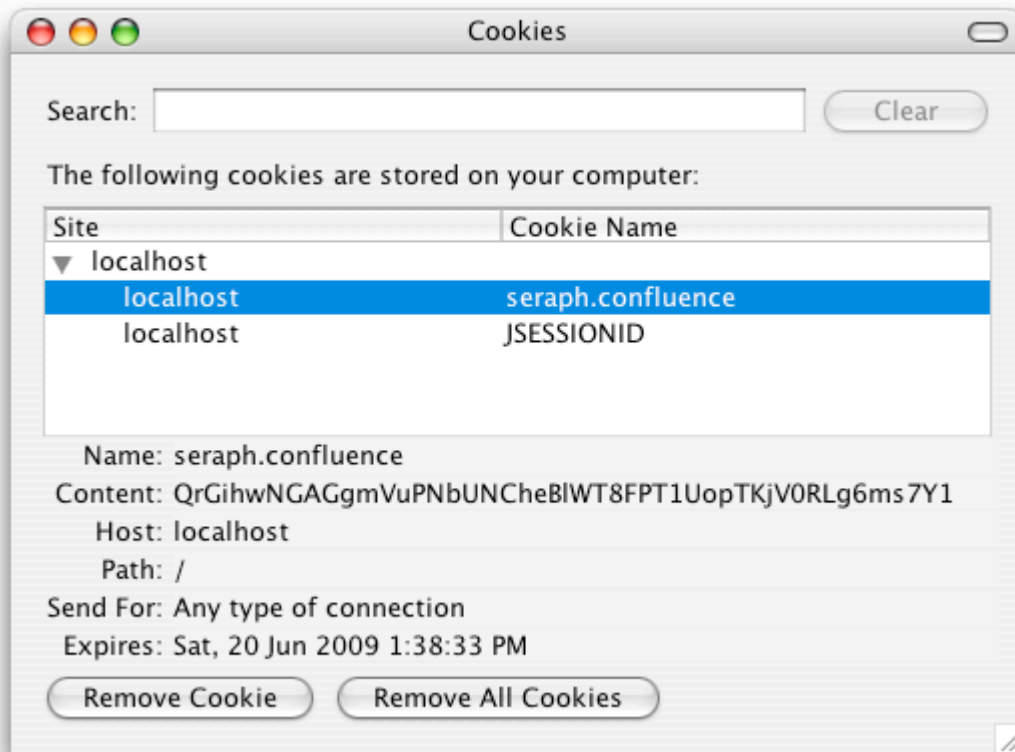
Safe Information Transit

The cookie information is always encoded by the server before it is given to a client. A cookie that has been tampered with will be considered to be not valid.

Session and Cookie Logic

Essentially, the cookie contains encrypted username and the user's password. To be more precise confluence uses PBE (password based encryption) with MD5 and DES, where password(also known as 'private key') is configurable in the `seraph.xml` file. The user's password in the cookie is necessary to ensure that the cookie is no longer valid if the user changes their password. The username must be retrievable by the server to identify the user solely from the cookie, which is what the 'Remember my login on this computer' feature does.

> ✅ The private key for confluence is stored at *confluence-install/confluence/WEB-INF/Classes/seraph.xml* under the
> *cookie.encoding* parameter. Please change this to something other than the default.

Is it Possible to Disable the 'Remember my login on this computer' Feature?

At the moment there is no available option for disabling "Remember My Login on this computer" feature via the Admin console. See the workaround here.

The user login *Auto Completion* functionality is a browser feature, and there is nothing Confluence can enable or disable.

RELATED TOPICS

Assigning Space Permissions

Confluence Cookies

Confluence Security

View File Macro

Revoking Space Permissions

Edit in Word Link Macro

Hiding the People Directory

Confluence Security Advisory 2006-01-23

Security Overview

HTML Macro

How to Hide the Referrer

Space Permissions Overview

Administrators Guide Home Confluence Documentation Home

Enabling or Disabling Public Signup

Enabling '**Public Signup**' allows users to sign themselves up to the site.

If you want to restrict your site to a particular set of users, you may want to disable 'Public Signup'. In this instance, administrators can add new users from the **Administration Console**.

To enable or disable public signup,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. This will display the '**General Configuration**' screen. Click '**Edit**'.
4. Beside '**Public Signup**', select '**On**' to enable Public Signup. Select '**Off**' to disable it.
5. Click '**Save**'.

RELATED TOPICS

Disabling the Built-In User Management
User Management
Security

Administrators Guide Home Confluence Documentation Home

Hiding External Links From Search Engines

Hiding external links from search engines helps to discourage spammers from posting links on your site. If you turn this option on, any URLs inserted in pages and comments will be given the 'nofollow' attribute, which prevents search engines from following them.

Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

To hide external links from search engines,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. This will display the '**General Configuration**' screen. Click 'Edit'.
4. Select 'On' beside '**Hide External Links From Search Engines**'.
5. '**Save**' your changes.

Background to the nofollow attribute
As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with some markup which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

RELATED TOPICS

Enabling or Disabling Public Signup

Adding SSL for Secure Logins and Page Security

Hiding external referrers

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Hiding the People Directory

The People Directory provides a list of all users in your Confluence system.

If you need to disable the People Directory set the following system properties on your application server command line:

- **To disable the People Directory for anonymous users,**

```
-Dconfluence.disable.peopledirectory.anonymous=true
```

- **To disable the People Directory entirely,**

```
-Dconfluence.disable.peopledirectory.all=true
```

To remove the link on the dashboard:

> ℹ️ This only applies to Confluence 2.5.2 to 2.9.x. Confluence 2.10.x or later only needs to configure system properties using the above.
> Edit the <confluence-install>/confluence/decorators/global.vmd:
> Comment out line 37:
>
> ```
> 1.<!--                      <img src="$req.contextPath/images/icons/people_directory_32.gif"
> align='absmiddle' height="32" width="32"> <b><a class="fontSizeDefault" href=
> "$req.contextPath/peopledirectory.action"> $action.getText("people.directory.title"
> )</a></b><span class="smalltext"> -    $action.getText("people.directory.description"
> )</span><br> -->
> ```

RELATED TOPICS

📄 Enabling or Disabling Public Signup

📄 Adding SSL for Secure Logins and Page Security

📄 Hiding external referrers

📄 Excluding external referrers

📄 Hiding External Links From Search Engines

📄 Hiding the People Directory

📄 User Email Visibility

📄 Anonymous Access to Remote API

📄 Ignoring External Referrers

📄 Spam Prevention via Captcha

Managing External Referrers

Administrators Guide Home Confluence Documentation Home

Managing External Referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

By default, external referrers for a page are listed under '**Hot Referrers**' on the '**Info**' screen of the page. (See *Screenshot 1* below.) Confluence shows a maximum of 10 referrers. If there are more than 10, confluence shows the 10 with the highest number of hits.

Note that you do *not* need to enable trackback in order to have external referrers enabled.

To manage your external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select the '**Manage Referrers**' option (See *Screenshot 2* below.).

The following actions will be available:

- Record or ignore all external referrers: By default, Confluence records the number of hits made to a page from the link on the external site. If you turn this option off, Confluence will not record the hits.
- Show or hide all external referrers: By default, Confluence lists the external referrers as '**Hot Referrers**' on the '**Info**' screen of a page, as shown below. If you turn this option off, external referrers will not be listed on the page.
- Specify which external referrers to exclude: You can decide which referrers you want to exclude from being displayed on your site.

Screenshot 1: Hot Referrers showing on a page's Info screen

Screenshot 2: Managing external referrers

Record External Referrers:	On | Off
Show Referrers in Page Info:	On | Off
Excluded External Referrer Prefixes:	[_____] [Add] ☐ Purge All http://www.google.com (del | purge) http://www.yahoo.com (del | purge)

RELATED TOPICS

📄 Enabling or Disabling Public Signup (Confluence Docs 3.1)

📄 Adding SSL for Secure Logins and Page Security (Confluence Docs 3.1)

📄 Hiding external referrers (Confluence Docs 3.1)

📄 Excluding external referrers (Confluence Docs 3.1)

📄 Hiding External Links From Search Engines (Confluence Docs 3.1)

📄 Hiding the People Directory (Confluence Docs 3.1)

📄 User Email Visibility (Confluence Docs 3.1)

📄 Anonymous Access to Remote API (Confluence Docs 3.1)

📄 Ignoring External Referrers (Confluence Docs 3.1)

📄 Spam Prevention via Captcha (Confluence Docs 3.1)

📄 Managing External Referrers (Confluence Docs 3.1)

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Excluding external referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

You can exclude external referrers to prevent them from being recorded or displayed anywhere on your site. Once you have specified your list of blocked URLs, any incoming links from URLs that match the list will no longer be recorded. Referrer URLs are blocked if they start with any of the URLs in the exclusion list. So http://evilspamsite.blogspot.com will also match http://evilspamsite.blogspot.com/nastypage.html

There are two instances where you may want to do this:

1. If you are running a Confluence installation that is open to public:
   In a site that is open to public, one unfortunate problem is that malicious sites can spam the display of a page's incoming links statistics. This is usually done to get the site's URL to appear in the sidebar. By adding these sites to the 'excluded referrers' list, you can prevent them from being listed on your site.
2. If Confluence is installed on a server with multiple domain names or IP addresses:
   Confluence will consider any URL originating from the domain name where Confluence is installed as an internal link. However, if Confluence is installed on a server with multiple domain names or IP addresses, you will need to add the other domain name prefixes to this list to let Confluence know that any links from these domains should not be considered external links.

ℹ️ You need to be a Confluence administrator and to know the URL of the site to add it to the excluded referrers list.

To add a URL to the excluded referrers list,

1. Go to the Confluence '**Administration Console**'. To do this:

   * Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Add the URL to the '**Excluded External Referrer Prefixes**' section.
   * You must include 'http://' at the front of the URL.
   * You can add more than one URL by putting each URL on a new line.

Screenshot: Excluding external referrers

Record External Referrers:	On | Off
Show Referrers in Page Info:	On | Off
Excluded External Referrer Prefixes:	[ ] Add ☐ Purge All http://www.google.com (del | purge) http://www.yahoo.com (del | purge)

RELATED TOPICS

Enabling or Disabling Public Signup

Adding SSL for Secure Logins and Page Security

Hiding external referrers

Excluding external referrers

Hiding External Links From Search Engines

Hiding the People Directory

User Email Visibility

Anonymous Access to Remote API

Ignoring External Referrers

Spam Prevention via Captcha

Managing External Referrers

Administrators Guide Home Confluence Documentation Home

Hiding external referrers

By default, Confluence lists the external referrers as '**Hot Referrers**' on the '**Info**' screen of a page. If you turn this option off, external referrers will not be listed on the page.

To hide external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Click '**Off**' beside '**Show Referrers in Page Info**'.

Screenshot: Managing external referrers

Record External Referrers:	On | Off
Show Referrers in Page Info:	On | Off
Excluded External Referrer Prefixes:	[ ] Add ☐ Purge All

RELATED TOPICS

📄 Enabling or Disabling Public Signup

📄 Adding SSL for Secure Logins and Page Security

📄 Hiding external referrers

📄 Excluding external referrers

📄 Hiding External Links From Search Engines

📄 Hiding the People Directory

📄 User Email Visibility

📄 Anonymous Access to Remote API

📄 Ignoring External Referrers

📄 Spam Prevention via Captcha

📄 Managing External Referrers

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Ignoring External Referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral. By default, Confluence records the number of hits made to a page from any link on an external site. If you turn this option off, Confluence will not record the hits.

To ignore external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Click '**Off**' beside '**Record External Referrers**'.

Screenshot: Managing external referrers

Record External Referrers:	On | Off
Show Referrers in Page Info:	On | Off
Excluded External Referrer Prefixes:	[_____] [Add] ☐ Purge All

RELATED TOPICS

📄 Enabling or Disabling Public Signup

📄 Adding SSL for Secure Logins and Page Security

📄 Hiding external referrers

📄 Excluding external referrers

📄 Hiding External Links From Search Engines

📄 Hiding the People Directory

📄 User Email Visibility

📄 Anonymous Access to Remote API

📄 Ignoring External Referrers

📄 Spam Prevention via Captcha

📄 Managing External Referrers

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

User Email Visibility

Confluence provides three options for email address privacy which can be configured by a Confluence administrator from the **Administration Console**:

- **Public**: email addresses are displayed publicly.

- **Masked**: email addresses are still displayed publicly, but masked in such a way to make it harder for spam-bots to harvest them.

- **Only visible to site administrators**: only Confluence administrators can see the email addresses. Note that, if you select this option, email addresses will not be available in the 'User Search' popup (e.g. when setting Page Restrictions).

To configure user email visibility,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. This will display the '**General Configuration**' screen. Click '**Edit**'.
4. Select one of the options next to '**User email visibility**': '**public**', '**masked**', or '**only visible to site administrators**'.
5. '**Save**' your changes.

Screenshot : email visibility

User email visibility:	○ public ○ masked (i.e. user at example dot com) ◉ only visible to site administrators

RELATED TOPICS

Enabling or Disabling Public Signup

Adding SSL for Secure Logins and Page Security

Hiding external referrers

Excluding external referrers

Hiding External Links From Search Engines

Hiding the People Directory

User Email Visibility

Anonymous Access to Remote API

Ignoring External Referrers

Spam Prevention via Captcha

Managing External Referrers

Administrators Guide Home   Confluence Documentation Home

Spam Prevention via Captcha

You need to be a Confluence administrator to enable Captcha.

If your Confluence site is open to the public you may find that automated spam is being added, in the form of comments or new pages.

You can configure Confluence to deter automated spam by asking users to prove that they are human before they are allowed to:

- Sign up for an account
- Add a comment
- Create a page
- Edit a page

Captcha is the technical term for a test that can distinguish a human being from an automated agent such as a web spider or robot.

When Captcha is switched on, users will need to recognise a distorted picture of a word, and must type the word into a text field. This is easy for humans to do, but very difficult for computers.

You can configure Confluence to enforce Captcha for certain types of users. You can exempt logged-in users (they will have completed a Captcha when they signed up), or members of particular groups.

By default, Captcha images will not be shown to logged-in users. Only anonymous users will have to perform the Captcha test when creating comments or editing pages.

To enable Captcha for Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Spam Prevention**' from the 'Configuration' menu on the left.
3. Turn on Captcha by clicking the '**ON**' link.
4. If you want to disable Captcha for certain groups:
   - Select '**No one**' if you want everyone to see Captchas.
   - Select '**Signed in users**' if you want only anonymous users to see Captchas.
   - If you want everyone to see Captchas except members of specific groups, select the '**Members of the following groups**' and enter the group names in the text box.
     You can click the magnifying-glass icon to search for groups. Search for all or part of a group name and click the '**Select Groups**' button to add one or more groups to the list.
   - To remove a group from the list, delete the group name.
5. Click the '**Save**' button.

Working with Confluence Logs

On this page:

Logs Location

This section describes Confluence's default logging behaviour, assuming that you have not changed the destination of the logs. So as to unify logging across different application servers, Confluence uses the **atlassian-confluence.log**, not the application server log, as it's primary log.

> Atlassian Support will almost always ask for the **atlassian-confluence.log** from the `confluence-home/logs` directory. The easiest way to find this location is to look for the "Confluence Home" setting from `Administration > System Information`. If you can't access `Administration > System Information`, check `<confluence-install>/confluence/WEB-INF/classes/confluence-init.properties` and look for the confluence.home setting, then find the logs in that directory.

For **Confluence 2.6.x and earlier**, the default behaviour is:

- For Confluence Standalone, log entries are written to `<confluence_install>/logs`. The main log file is called `atlassian-confluence.log`.
- For Confluence EAR/WAR, log entries are written to the application server logs, i.e. the default log files of the application container.

For **Confluence 2.7.x and later**, both Standalone and EAR/WAR distributions follow the same default behaviour:

- When you start Confluence, log entries will be sent to the application server logs until Confluence has completed its initial bootstrap. Any log entries written to the console will be repeated into the `<confluence-home>` log described below.
- Once the initial startup sequence is complete, all logging will be to `<confluence-home>/logs/atlassian-confluence.log`. For example: `c:/confluence/data/logs/atlassian-confluence.log`.

  Note that the default location is now the Confluence **home directory** instead of the **application server's log file**. The home directory is specified in <confluence-install>/confluence/WEB-INF/classes/confluence-init.properties.

Background

Confluence uses Apache's log4j logging service. This allows a developer or administrator to control the logging behavior and the log output file by editing a configuration file, without touching the application binary. There are six known log4j logging levels.

Finding the Log Configuration File

Confluence's logging behaviour is defined in the following properties file:
`<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`

This file is a standard log4j configuration file, as described in the Apache log4j documentation.

Changing the Destination of the Log Files

Terminology: In log4j, an output destination is called an 'appender'.

To change the destination of the log files, you need to stop Confluence and then change the settings in the '**Logging Location and Appender**' section of the `log4j.properties` file. The location of this file is described above.

In the standard properties file supplied with **Confluence 2.7 and later**, you will find entries for two appenders:

- `com.atlassian.confluence.logging.ConfluenceHomeLogAppender` – This is a custom appender which controls the default logging destination described above. This appender allows the following settings:
    - MaxFileSize
    - MaxBackupIndex
- `org.apache.log4j.RollingFileAppender` – If you want to log to a different location, uncomment the `RollingFileAppender` line and change the destination file in the line below it. Comment out the previous lines referring to the `ConfluenceHomeLogAppender`.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the log4j documentation.

Changing the Logging Levels

See Configuring Logging for instructions on how to change the logging configuration of Confluence.

Using Some Specific Confluence Logging Options

This section contains some pointers to specific log configurations you may need.

Log the Details of SQL Requests made to the Database

You may want to increase Confluence's logging so that it records individual SQL requests sent to the database. This is useful for troubleshooting specific problems.

You can enable detailed SQL logging in two ways:

- At runtime – see instructions above.
- Via the logging properties file – see the detailed instructions.

Log the Details of Users Viewing/Accessing each Confluence Page

You can configure the log to show which users are accessing which pages in Confluence. This can only be done via the logging properties file – see the detailed instructions.

Where are my Thread Dumps?

Thread dumps are logged to the application server log file.

RELATED TOPICS

- Important Directories and Files
- Enabling detailed SQL logging
- Enabling user access logging
- Generating a Thread Dump
- Enabling Page Request Profiling

log4j Logging Levels

Logging Levels

- **DEBUG** - designates fine-grained informational events that are most useful to debug an application (*what is going on*)

- **INFO** - announcements about the normal operation of the system - scheduled jobs running, services starting and stopping, user-triggered processes and actions

- **WARN** - any condition that, while not an error in itself, may indicate that the system is running sub-optimally

- **ERROR** - a condition that indicates something has gone wrong with the system

- **FATAL** - a condition that indicates something has gone wrong so badly that the system can not recover

- **TRACE** - n/a within confluence

> There are two ways to modify the logging levels, as described in Working with Confluence Logs.
>
> 1. Modifying the runtime log levels via the **Administration Console**.
> 2. Manually modifying the `<Confluence-Install>\confluence\WEB-INF\classes\log4j.properties` file.

Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

Reference : log4j manual

Configuring OAuth

On this page:

What is OAuth?

OAuth is a protocol that allows a web application to share a finite set of its private resources and data (through gadgets, for example) with another OAuth-compliant application. These applications could be a Confluence, a JIRA or a website such as iGoogle.

Using OAuth, you can access data within a Confluence installation externally, such as on a JIRA site's dashboard, another Confluence site's page, or a website like iGoogle. This would be done via a gadget supplied by Confluence. While some data in Confluence may be accessible anonymously on the external application, other data may only be available to a specific user account within the Confluence installation. OAuth provides this capability.

The key security advantage of OAuth is that a web application's resources, which are only available to a specific user, can be shared without the web application having to hand out the user's authentication details. Instead, access to these private resources is handled via an '**access token**'. Access tokens typically define what web application resources (usually based on user account privileges) can be accessed by another application and the duration of this access. However, access tokens are dissociated from a user's authentication details, since authentication to gain access to these resources is handled separately.

> In OAuth terminology, an application that shares its resources is known as a **service provider** and an application that accesses a service provider's resources is known as a **consumer**.

For more information about OAuth, please refer to the OAuth protocol workflow section of our Gadgets and Dashboards documentation. It is important to understand this workflow first before establishing OAuth relationships between your Confluence installation and other external web applications (either Atlassian or non-Atlassian ones).

> **Important information about establishing OAuth relationships for gadgets**
>
> If you wish to use a gadget served by any Atlassian application and access data restricted to a user account on that application, then an OAuth relationship between the service provider and consumer application *should be established first*.
>
> Alternatively, if the gadget is served by an Atlassian application that supports Atlassian's Trusted Applications feature (for example, JIRA, Confluence or Bamboo), you can establish a Trusted Applications relationship instead of an OAuth one. Bear in mind that in Trusted Application relationships, you can only access data restricted to a user account on the service provider if:
>
> 1. The usernames of user accounts on the service provider and consumer applications match.
> 2. The user has logged in to the consumer application.
>
> OAuth relationships provides an individual with access to restricted data on the service provider if their usernames on the service provider and consumer applications, differ since authentication is part of the OAuth protocol workflow.
>
> Not all external gadgets used in Confluence require the establishment of an OAuth relationship. If the gadget does not need to access restricted resources on the service provider, then there should be no need to establish an OAuth relationship.
>
> The instructions in this section provide information on how to establish an OAuth relationship between your Confluence site and another web application's site. This could apply to instances where Confluence acts as either the consumer or service provider in the relationship.

Accessing and Using Confluence's OAuth Administration Page

Confluence's OAuth Administration section, which handles the establishment of OAuth relationships between consumer and service provider web applications, is found in the Administration Console area of Confluence.

To access Confluence's OAuth Administration page,

1. Go to the Confluence '**Administration Console**'. To do this:

    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**OAuth**' from the 'Administration' section on the left navigation panel.

To use the OAuth administration page,

- On the 'OAuth Administration' page:
  - Click the '**Consumers**' tab to configure consumer applications that will be accessing the resources of your Atlassian application. Refer to Configuring OAuth Consumers for more information.
  - Click the '**Consumer Info**' tab to view or edit your Atlassian application's Consumer information. Refer to Configuring OAuth Consumer Information for more information.
  - Click the '**Service Providers**' tab to configure service providers whose resources your Atlassian application will be consuming. Refer to Configuring OAuth Service Providers for more information.

ℹ In the procedure above (and referred documentation), 'your Atlassian application' refers to your Confluence installation.

User Management

- Confluence User Management
  - Searching For and Managing Users
  - Adding a Group
  - Adding a New User
  - Adding or Removing Users in Groups
  - Changing Usernames
  - Editing User Details
  - Global Groups Overview
  - Global Permissions Overview
  - Migrating to new User Management
  - Removing a Group
  - Removing a User
  - Setting up Anonymous Access
  - Viewing members of a group
  - How to Improve User Search Performance — If your Confluence instance contains thousands of user accounts and you are experiencing performance issues when searching for users, the following migration guide is for you.
  - Restoring Passwords To Recover Admin User Rights
- Integrating with Crowd
- JIRA User Management
  - Delegate user management to use JIRA logins
    - Revert from JIRA to internal user management
- LDAP User Management
  - Add LDAP Integration
    - Automatically Add LDAP users to the confluence-users Group
    - Customising atlassian-user.xml
    - Migrate to LDAP User Management From OsUser
  - Add LDAP Integration For User Authentication Only
  - atlassian-user.xml reference — || XML tag || Default value || Description |
  - Changes in osuser.xml from 1.0.3a to 1.1.x
  - Configuring multiple LDAP repositories
  - Connect to LDAP, JIRA or Other Services Via SSL
  - Disabling the Built-In User Management
  - Legacy User Management Documentation
    - LDAP Authentication with OSUser
  - Troubleshooting LDAP User Management
  - Troubleshooting the "Not Permitted" Screen under LDAP Integration
- Migrating users from Confluence to JIRA — There is currently no way to delegate user management from JIRA to Confluence. So, if you are in a situation where your users are defined in Confluence and would like to take advantage of Confluence's ability to use JIRA user management, you will need to transfer all of your existing Confluence users into JIRA. You can do this manually, or if you have a large number of users, you can use the attached XML-RPC script.
- Requesting External User Management Support
  - Paddle
- Understanding User Management in Confluence
- User Management Frequently Asked Questions

Confluence User Management

- Searching For and Managing Users
- Adding a Group
- Adding a New User
- Adding or Removing Users in Groups
- Changing Usernames

- Editing User Details
- Global Groups Overview
- Global Permissions Overview
- Migrating to new User Management
- Removing a Group
- Removing a User
- Setting up Anonymous Access
- Viewing members of a group
- How to Improve User Search Performance
- Restoring Passwords To Recover Admin User Rights

Searching For and Managing Users

If you are a Confluence Administrator, you can add users, assign them to groups and edit their user details.

On this page:

- Accessing the User Management Screen
- Listing All Users
- Using the Simple User Search
- Using the Advanced User Search

Accessing the User Management Screen

To search for and manage users,

1. Go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
     - Select '**Manage Users**' in the left-hand panel.
     - The '**Manage Users**' screen appears, as shown below. You can now choose to list all users or you can search for a specific user.

Screenshot: Manage users
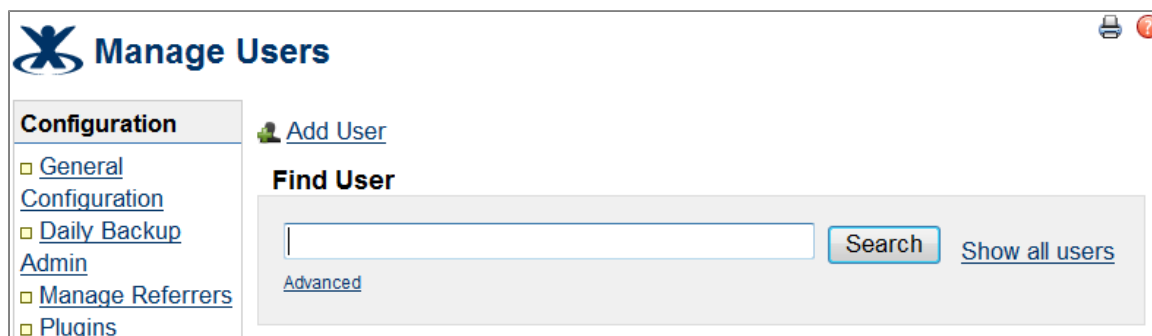
Listing All Users

To list all users,

1. Open the '**Manage Users**' screen as described above.
2. Click the '**Show all users**' link. All members of the confluence-users group are listed in alphabetical order, by username. If there are more users than can fit on one page, the results will be divided into multiple pages.
3. To move to another page of results, click the numbered links, '**Next**' or '**Previous**' near the top or bottom of the page.
4. To specify how many results should be shown per page, click a number '**10**', '**20**', '**50**' or '**100**' near the top of the page.

Screenshot: List all users

Using the Simple User Search

> ⓘ **Crowd and the User Search**
> If you are using Atlassian's Crowd for user management, you will need Crowd 1.5.1 or later to use the 'Simple' option in the user search. If your version of Crowd does not support the simple user search, you will see only the 'Advanced' search form.

To search for a specific user via the simple user search,

1. Open the '**Manage Users**' screen as described above.
2. If the '**Simple**' link is showing, click it. (If you see the 'Advanced' link and no 'Simple' link, then you're fine. The simple search is already active.)
3. The simple user search screen will appear, as shown below.
4. Type some information about the user into the 'Search' textbox. You can type all or part of their username, full name or email address.
5. Click the '**Search**' button.
6. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Screenshot: Simple user search

Using the Advanced User Search

The advanced user search allows you to specify the field in which your search term appears, i.e. username, full name or email address. You may find this useful if you need to limit the number of users appearing in the search results.

To search via the advanced user search,

1. Open the '**Manage Users**' screen as described above.
2. If the '**Advanced**' link is showing, click it. (If you see the 'Simple' link and no 'Advanced' link, then you're fine. The advanced search is already active.)
3. The advanced user search screen will appear, as shown below.
4. Complete one or more of the following fields:
   - **User Name** — Enter all or part of the person's username i.e. their login id, e.g. 'joe', or 'bloggs'.
   - **Full Name** — Enter all or part of the person's name, e.g. 'joe bloggs', or 'bloggs', or 'joe'.
   - **E-Mail** — Enter all or part of the person's email address, e.g. 'acme'
5. Click the '**Search**' button.
6. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Screenshot: Advanced user search

RELATED TOPICS

Showing first 15 of 16 results

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Adding a Group

To add a new group,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left panel.
3. Enter a name for your group in the '**Add Group**' input field and click '**Save**'.

You are now ready to start adding users to the group.

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Searching For and Managing Users

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Adding a New User

There are two ways a new user can be added to Confluence:

Public Signup: Enabling public signup from the **Administration Console** allows users to sign themselves up to the site.

By Confluence Administrators : If you want to restrict your site to a select group of users, you may want to disable 'Public Signup'. In this instance, administrators with Confluence Administrator or System Administrator permissions can add new users from the **Administration Console**.

To add a new user to Confluence from the Administration Console,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Users**' in the left-hand panel.
3. Click the link '**Add new user**' at the top of the page.
4. In the form displayed, enter the user's details: username, password, name and email address.
5. Click '**Create**' to add the user.

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

Changing Usernames

Setting up Anonymous Access

Adding a Group

Security Overview

Troubleshooting LDAP User Management

Adding a New User

Global Permissions Overview

Disabling the Built-In User Management

Showing first 15 of 16 results

Administrators Guide Home Confluence Documentation Home

Adding or Removing Users in Groups

If you are a Confluence Administrator, you can add users and groups, and assign users to groups in order to determine their permissions.

This page tells you how to add a user to a group or remove a user from a group. For an overview of users and groups, please refer to Users and Groups and Confluence User Management.

You can edit group membership in two places:

- From the group management screen.
- From the user management screen for a particular user.

Both methods are described below.

On this page:

- Adding and Removing Members via the Group Management Screen
- Editing Group Membership from the User Management Screen
- A Note about External User Directories

Adding and Removing Members via the Group Management Screen

This is the recommended method, available in **Confluence 2.10** and later. It allows you to manage the group membership for a number of users at the same time.

To add members to a group,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left-hand panel.
3. The '**Manage Groups**' screen appears, showing a list of groups. Select the group to which you want to add users.
4. The '**Group Members**' screen appears, showing the users who belong to the selected group. (See screenshot below.) Click the '**Add Members**' link.
5. The '**Add Members**' screen appears, as shown below. Type in the usernames of the people you want to add to the group. You can also search for and select users by clicking the 🔍 icon, as described in Searching for Users.
6. When you have added the required username(s), click the '**Add**' button to add the member(s) to the group.

To remove members from a group,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left-hand panel.
3. The '**Manage Groups**' screen appears, showing a list of groups. Select the group from which you want to remove the user.
4. The '**Group Members**' screen appears, showing the users who belong to the selected group. (See screenshot below.) Click the 'Remove user from group' 👤 icon next to the user whose group membership you want to remove.

Screenshot: Group Members

Screenshot: Add Members

Editing Group Membership from the User Management Screen

You can update a user's group membership from the user management screen. This functionality allows you to update one user at a time.

To add a user to a group or remove a user from a group,

1. Go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
     - Select '**Manage Users**' in the left-hand panel.
     - The '**Manage Users**' screen appears, as shown below. You can now choose to '**Show all users**' or you can search for a specific user by entering all or part of the person's username, full name or email address. (For more details about the user search, see Searching For and Managing Users.)
     - Click the link on the username you want to edit.
2. Now you should be able to see the user's current details, with links allowing you to edit the user's details and groups. See the screenshot showing a user's details below.

3. Click '**Edit Groups**'. This will display two lists of groups, as shown in the screenshot below. Update the user's group membership as follows:
   - '**Not a member of groups**' — This box shows all groups to which the user does not belong. To add the user to a group, select a group and click '**Join**'. Hold the Ctrl key down and click to select more than one group.
   - '**Member of groups**' — This box shows all groups to which the user belongs. Select a group and click '**Leave**' to remove the user from the group.

Screenshot: Manage users

Screenshot: User details

Screenshot: Editing a user's groups

A Note about External User Directories

If you are using external user management, you cannot use the Confluence administration screens to add or remove users in external groups, as Confluence access to LDAP and JIRA groups is read only.

If you have Confluence integrated with Crowd and external user management turned **off** in Confluence, then your group membership changes in Confluence will be passed through to the external directory (e.g. LDAP) managed by Crowd (assuming that Crowd has read-write access to the directory).

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

Changing Usernames

Setting up Anonymous Access

Adding a Group

Security Overview

Troubleshooting LDAP User Management

Adding a New User

Global Permissions Overview

Disabling the Built-In User Management

Removing a User

Administrators Guide Home   Confluence Documentation Home

Changing Usernames

A **username** is the name used to log into Confluence, eg. `jsmith`.

> Currently, there is no straightforward method for changing a username and its associated content, to that of another user. The only practicable method currently available is to execute direct SQL queries on your database. There is a feature request to facilitate this process via a web interface and you can vote for it to improve its chances of being implemented. Be aware, however, that no matter what method you use to change usernames in Confluence, there is **no support** provided for this process. The instructions below provide suggested guidelines on how to change a username via SQL queries, although this may vary depending on your database.

Instructions For Changing Usernames

> ⚠ The following SQL commands are only tested for MYSQL and POSTGRES Databases. If you have any other database please contact your DBA to determine the equivalent queries.

Usernames can only be changed through direct update to the Confluence database.

1. If you have a database administrator, request that they approve the database-related steps described below

2. If you are using JIRA user management, Revert from JIRA To Internal User Management

3. Backup Confluence

4. Creating a `usermigration` table:

```
1. create table usermigration
2. (
3. oldusername varchar,
4. newusername varchar
5. )
```

5. Usernames that will be changed must be placed in the `usermigration` table with their current and planned usernames:

```
1. insert into usermigration (oldusername, newusername)
2. values ('oldusername', 'newusername');
```

6. If your DB administration tool does not support multiple SQL queries, these must be entered individually.

```
001. update attachments
002. set creator = newusername from usermigration u
003. where creator = u.oldusername;
004.
005. update attachments
006. set lastmodifier = newusername from usermigration u
007. where lastmodifier = u.oldusername;
008.
009. update content
010. set lastmodifier = newusername from usermigration u
011. where lastmodifier = u.oldusername;
012.
013. update content
014. set creator = newusername from usermigration u
015. where creator = u.oldusername;
016.
017. update content
018. set username = newusername from usermigration u
019. where username = u.oldusername;
020.
021. update extrnlnks
022. set creator = newusername from usermigration u
023. where creator = u.oldusername;
024.
025. update extrnlnks
026. set lastmodifier = newusername from usermigration u
027. where lastmodifier = u.oldusername;
028.
029. update label
030. set owner = newusername from usermigration u
031. where owner = u.oldusername;
032.
033. update content_label
034. set owner = newusername from usermigration u
035. where owner = u.oldusername;
036.
037. update links
038. set lastmodifier = newusername from usermigration u
039. where lastmodifier = u.oldusername;
040.
041. update links
```

```
042.set creator = newusername from usermigration u
043.where creator = u.oldusername;
044.
045.update notifications
046.set lastmodifier = newusername from usermigration u
047.where lastmodifier = u.oldusername;
048.
049.update notifications
050.set creator = newusername from usermigration u
051.where creator = u.oldusername;
052.
053.update pagetemplates
054.set lastmodifier = newusername from usermigration u
055.where lastmodifier = u.oldusername;
056.
057.update pagetemplates
058.set creator = newusername from usermigration u
059.where creator = u.oldusername;
060.
061.update spaces
062.set creator = newusername from usermigration u
063.where creator = u.oldusername;
064.
065.update spaces
066.set lastmodifier = newusername from usermigration u
067.where lastmodifier = u.oldusername;
068.
069.update spacepermissions
070.set permusername = newusername from usermigration u
071.where permusername = u.oldusername;
072.
073.update spacepermissions
074.set creator = newusername from usermigration u
075.where creator = u.oldusername;
076.
077.update spacepermissions
078.set lastmodifier = newusername from usermigration u
079.where lastmodifier = u.oldusername;
080.
081.update contentlock
082.set creator = newusername from usermigration u
083.where creator = u.oldusername;
084.
085.update contentlock
086.set lastmodifier = newusername from usermigration u
087.where lastmodifier = u.oldusername;
088.
089.update os_user
090.set username = newusername from usermigration u
091.where username = u.oldusername;
092.
093.update trackbacklinks
094.set creator = newusername from usermigration u
095.where creator = u.oldusername;
096.
097.update trackbacklinks
098.set lastmodifier = newusername from usermigration u
099.where lastmodifier = u.oldusername;
```

7. If using Confluence 2.1 or newer, run the following command:

```
1.update users
2.set name = newusername from usermigration u
3.where name = u.oldusername;
```

8. Reassign personal spaces and content associated with the old username to the new username. The tilda (~) is required as it is prepended to the space key of all personal spaces.

```
01.update spaces
02.set spacekey = '~newusername'
03.where spacekey = '~oldusername';
04.
05.update content_label
06.set spacekey = '~newusername'
07.where spacekey = '~oldusername';
08.
09.update bandana
10.set bandanacontext = '~newusername'
11.where bandanacontext = '~oldusername';
```

9. Each username is associated with a full name. For example, username 'jsmith' may have a full name of 'John M Smith'. If this

fullname needs to be changed, modify the fullname in the `users` or `os_user` table.

All old usernames in Confluence should now be replaced with the new usernames from the `usermigration` table.

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

Changing Usernames

Setting up Anonymous Access

Adding a Group

Security Overview

Troubleshooting LDAP User Management

Adding a New User

Global Permissions Overview

Disabling the Built-In User Management

Showing first 15 of 16 results

Administrators Guide Home   Confluence Documentation Home

Editing User Details

To update a user's details,

1. First, go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
       - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

       - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
       - Select the link '**Manage Users**' in the left-hand panel.
       - Locate the user by doing a search on the username or the groups to which they belong.
       - Click the user link.
2. Now you should be able to see the user's current details and links allowing you to edit them.
   - **View Profile** — View the user's profile.
   - **Edit Groups** — Add or remove this user from a group.
   - **Edit Details** — Edit the user's name and email address. Changing a user's username is not supported through the application, see Changing Usernames for other solutions.
   - **Set Password** — Edit the user's password details.
   - **Deactivate** — Deactivation is an alternative to removing a user. It is not possible to remove a user if the user is responsible for content on the site. The reason is that Confluence will need the user information to maintain a history of the pages they have updated. In this case, you can deactivate the user so that they can no longer log in to Confluence.
   - **Remove** — You can remove a user permanently if the user has not added or edited any content on the site.

⚠ **Deactivating users (unavailable)** : This functionality was removed in Confluence 2.1.x. See Removing a User for more information.

Screenshot : User Details

User:	testfoo2@atlassian.com
Full Name:	test foo 2
Email:	testfoo2@atlassian.com
Groups:	☐ confluence-users

View Profile | **Edit Groups** | **Edit Details** | **Set Password** | **Deactivate** | **Remove**

RELATED TOPICS

Adding or Removing Users in Groups

Searching For and Managing Users

Editing User Details

Setting up Anonymous Access

Adding a New User

Global Permissions Overview

Spam Prevention via Captcha

Removing a User

Global Groups Overview

There are two special default groups in Confluence:

1. **confluence-administrators**: This is a group of 'super-users' who can access the '**Administration Console**' and perform site-wide administration. Members of this group can also see all spaces in the Confluence instance.

   Any user who is a member of this group has site-wide administration powers, regardless of any other setting. The settings on the Global Permissions screen do **not** affect the powers allowed to members of this group.

   > There is an outstanding request to remove the 'confluence-administrators' group from a future version of Confluence (see CONF-4616).

2. **confluence-users**: This is the default group for all new users. Permissions you assign to this group will be assigned to all newly signed-up users of Confluence.

Other user groups : A Confluence administrator can also group users together into user groups for more convenient administration. Once created, groups become available at the space and page levels to allow for flexible access control. A user in one of these groups will automatically be granted all permissions granted to the group.

Anonymous users : Confluence treats all users who do not log in when they access Confluence as being 'anonymous'. You can grant anonymous '**Use Confluence**' permission via the Global Permissions screen. This will allow non-registered users to access pages and spaces in Confluence. A space administrator can then further control anonymous access per space via the space permissions.

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

📄 Changing Usernames

📄 Setting up Anonymous Access

📄 Adding a Group

Showing first 10 of 16 results

🏠 Administrators Guide Home   🏠 Confluence Documentation Home

Global Permissions Overview

Permissions determine the actions which a user is allowed to perform within Confluence. Global permissions are one of the levels of permission provided by Confluence.

In order to assign these permissions, you must already have the global '**Confluence Administrator**' or '**System Administrator**' permission (described below). You can then assign global permissions to groups, individual users and anonymous users. Further permissions are granted from the space administration screens.

On this page:

- Overview of the Global Permissions
- Comparing the System Administrator with the Confluence Administrator Permission
- Comparing the Administrator Permissions with the confluence-administrators Group
- Updating Global Permissions

Overview of the Global Permissions

Global permissions control access across the whole Confluence site. Here is a list:

Global Permission	Description
Can Use	This is the most basic permission that allows users to access the site. ℹ️ Users with this permission count towards the number of users allowed by your license. See the information on removing/deactivating users.
Attach Files to User Profile	This allows the user to upload files to be stored in their user profile. ℹ️ This feature was made obsolete by the introduction of personal spaces in Confluence 2.2. Hence, this permission is no longer relevant. Attachments can be accessed from a user profile view (for example, an image within the 'About Me' field of a profile view) by attaching these files to a page within that user's personal space and referencing them using appropriate wiki markup code.
Update User Status	This allows the user to update their user status message, which can be seen on the user's profile, pages in their personal space and on various activity streams accessible to other Confluence users.
Personal Space	This permission allows the user to create a personal space.
Create Space(s)	This permission allows users to create new spaces within your Confluence site. When a space is created, the creator automatically has the '**Admin**' permission for that space and can perform space-wide administrative functions.
Confluence Administrator	This permission allows users to access the **Administration Console** that controls site-wide administrative functions. Users with this permission can perform most, but not all, of the Confluence administrative functions. See the comparison of 'System Administrator' and 'Confluence Administrator' below.
System Administrator	This permission allows users to access the **Administration Console** that controls site-wide administrative functions. Users with this permission can perform all the Confluence administrative functions, including the ones which the 'Confluence Administrator' permission does not allow. Users with this permission are listed on the '**Site Administrators**' page that is linked from '**Contact Administrators**' in the footer throughout the Confluence site. See the comparison of 'System Administrator' and 'Confluence Administrator' below. Refer also to the note about the 'confluence-administrators' group below.

> ℹ️ **The first system administrator is defined during installation**
> During the initial configuration of Confluence, the Setup Wizard asks for the username of the System Administrator. This user will have the 'System Administrator' permission and will be a member of the 'confluence-administrators' group.

Comparing the System Administrator with the Confluence Administrator Permission

New with **Confluence 2.7 and later** comes the ability to have two levels of administrator in Confluence:

- **System Administrator** – Users with this permission can perform all the Confluence administrative functions, including the ones

which the 'Confluence Administrator' permission does not allow.
- **Confluence Administrator** – Users with this permission can perform most, but not all, of the Confluence administrative functions.

Tip: The two-tier administration is useful when you want to delegate some administrator privileges to project managers or team leaders. You can give 'Confluence Administrator' permission to users who should be able to perform most administrative functions, but should not be able to perform functions that can compromise the security of the Confluence system.

The following functions are excluded from the '**Confluence Administrator**' permission:

Administration Screen	Excluded Function
General Configuration	The following functionality is disallowed: • Server Base URL • Remote API plugin • External user management • Public Signup
Daily Backup Admin	This function is disallowed entirely.
Plugins	This function is disallowed entirely.
Plugin Repository	This function is disallowed entirely.
Mail Servers	This function is disallowed entirely.
User Macros	This function is disallowed entirely.
Attachment Storage	This function is disallowed entirely.
Layouts	This function is disallowed entirely.
Custom HTML	This function is disallowed entirely.
Backup & Restore	This function is disallowed entirely.
SnipSnap Import	This function is disallowed entirely.
Logging and Profiling	This function is disallowed entirely.
Cluster Configuration	This function is disallowed entirely.

Comparing the Administrator Permissions with the confluence-administrators Group

The '**confluence-administrators**' group defines a set of 'super-users' who can access the Administration Console and perform site-wide administration. Members of this group can also see all pages and spaces in the Confluence instance except pages for which they are excluded by page restrictions (restrictions can be removed by members of the confluence-administrators group in the Space Admin screen if need be). The settings on the 'Global Permissions' screen do not affect the powers allowed to members of this group.

Granting the '**System Administrator**' or '**Confluence Administrator**' permission to a user will *not* automatically grant the user access to all spaces in the site. These permissions will only give access to the Administration Console.

Be aware, however, that users with 'System Administrator' can add themselves to the 'confluence-administrators' group and become a super-user.

> ⚠ **Confluence Administrator permission and confluence-administrators group are not related**
> Going by the names, you would think the 'confluence-administrators' group and the 'Confluence Administrator' permission are related – but they are not. To resolve confusion, we want to make explicit that granting a user or group 'Confluence Administrator' permission is **not** the same as granting them membership to the 'confluence-administrators' group. Granting the 'Confluence Administrator' permission enables access to only a subset of the administrative functions. Granting membership to the 'confluence-administrators' group, on the other hand, gives complete access.

⚠ There is an outstanding request to remove the 'confluence-administrators' group from a future version of Confluence (see CONF-4616).

Read more about global groups.

Updating Global Permissions

To edit the global permissions for a group or user,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the '**Security**' section of the left-hand panel.
3. The '**View Global Permissions**' screen appears. Click the '**Edit Permissions**' button.
4. The '**Edit Global Permissions**' screen appears, as shown below. Add or edit group and user permissions as follows:
   - To add permissions for a group:
     a. First add the group to Confluence, if you have not already done so.
     b. Now on the 'Edit Global Permissions' screen, enter the group name in the text box labelled '**Grant browse permission to**' in the '**Groups**' section. You can click the magnifying glass to search for the group name.
     c. Click the '**Add**' button.
     d. The group will appear in the list and you can now edit its permissions.
   - To add permissions for a specific user:
     a. First add the user to Confluence, if you have not already done so.
     b. Consider adding the user to a group and then assigning the permissions to the group, as described above, instead of assigning permissions to the specific user.
     c. To assign permissions to a specific user on the 'Edit Global Permissions' screen, enter the username in the text box labelled '**Grant browse permission to**' in the '**Individual Users**' section. You can click the magnifying glass to search for the username.
     d. Click the '**Add**' button.
     e. The username will appear in the list and you can now edit its permissions.
   - To add or edit the permissions for a user or group:
     - Select the check box under the relevant permission and next to the relevant user/group. A tick in the check box indicates that the permission is granted. Click again to clear the check box and deny the permission.
   - To allow anonymous access to your Confluence site, select the 'Use Confluence' and 'View User Profile' options in the 'Anonymous Access' section.

     ⓘ For more information about these permissions, refer to Setting up Anonymous Access.
   - Click the '**Save All**' button to save your changes.

Screenshot: Editing global permissions

Groups

These are the global permissions currently assigned to groups.

		Attach Files to User Profile	Update User Status	Personal Space	Create Space(s)	Confluence Administrator	System Administrator
👥 confluence-administrators	☑ can use	☐	☐	☑	☑	☑	☑
👥 confluence-users	☑ can use	☐	☐	☑	☐	☐	☐

Grant browse permission to [ ] 🔍 [Add]

Individual Users

These are the global permissions currently assigned to individual users.

No users currently have global permissions.

Grant browse permission to [ ] 🔍 [Add]

Anonymous Access

When a user is using Confluence while not logged in, they are using it anonymously.
For example: Enabling anonymous 'USE' permission, allows non-logged-in users to browse pages and spaces in Confluence.

	Use Confluence	View User Profiles
👥 Anonymous	☐ can use	☐

[Save All] [Cancel]

> **About some error messages you may see**
> In **Confluence 2.7.2** and later, Confluence will let you know if there is a problem with some permissions. In rare situations, you may see the following error messages below a permission:
>
> - '**User/Group not found**' — This message may appear if your LDAP repository is unavailable, or if the user/group has been deleted after the permission was created.
> - '**Case incorrect. Correct case is: xxxxxx**' — This message may appear if the upper/lower case in the permission does not match the case of the username or group name. If you see a number of occurrences of this message, you should consider running the routine supplied to fix the problem.

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

Changing Usernames

Setting up Anonymous Access

Adding a Group

Security Overview

Troubleshooting LDAP User Management

Adding a New User

Global Permissions Overview

Disabling the Built-In User Management

Showing first 15 of 16 results   Administrators Guide Home   Confluence Documentation Home

Migrating to new User Management

In order to support advanced forms of user management, Confluence now uses the AtlassianUser management framework, which can store users in the database (through Hibernate) instead of delegating the user management to OSuser.

> **Manual migration no longer required from Confluence 2.7.0**
> This page describes how to perform a manual migration of your users from OSuser to AtlassianUser. For **Confluence 2.7.0 and later**, there is no need to perform a manual migration of your users to the AtlassianUser framework. If you are installing Confluence 2.7.0 or later for the first time, you will automatically receive the AtlassianUser framework. If you are upgrading from an earlier version to Confluence 2.7.0 or later and have not changed the default user management configuration, your users will be automatically migrated. Refer to the details in the Confluence 2.7 Upgrade Guide.

> ⚠ If going from OsUser to LDAP, use Migrate to LDAP User Management From OsUser rather than this page. This page is intended for upgrading internal user management repositories.

For **Confluence 2.6.x** and earlier, Confluence delegates user management to OSuser by default. However, you may wish to migrate your users away from OSuser for one or more of the following reasons:

- In preparation to use external user management.
- To take advantage of the native AtlassianUser's much more efficient searching and user administration.
- Because you are having problems with OSuser, such as CONF-5218.

⚠ The migration instructions below are valid for Confluence version 2.2 and later.

Step 1 - Upgrade Confluence

Please check that you are running the latest version of Confluence. If not, we strongly recommend that you consider upgrading Confluence according to this guide. Confirm that you have upgraded successfully before trying to add LDAP to the new version.

Step 2 - Confluence User Migration

1. Find your Confluence base URL. To check this from Confluence, go to `Administration > General Configuration > Base Url`. Record this for later in the process.

2. Make a backup of your:
   - database
   - Confluence home directory
   - `confluence/WEB-INF/classes/atlassian-user.xml` (only if you have made changes)

   ⚠ If you do not create a backup, you cannot roll back to the old version if the migration is unsuccessful.
3. Download ldap_hibernate_osuser_atlassian-user.xml, rename it to `atlassian-user.xml` and copy it to your `confluence/WEB-INF/classes` directory. (You can overwrite the one that is there).
   If you have already set up LDAP in your `osuser.xml` file and wish to migrate to atlassian-user LDAP, then you need to uncomment the ldap section, and fill in the correct details (as described in Customising atlassian-user.xml). This will prevent users in your osuser table that exist in LDAP from being migrated over. If you haven't already set up LDAP in `osuser.xml` please do NOT uncomment the ldap section.

4. Restart Confluence.
5. Log in as a System Administrator, copy the address `http://<BASEURL>/<contextpath>/admin/osuser2atluser.jsp` and paste it into your browser's address bar. Change <BASEURL> to your actual base URL and <contextpath> to your context path (usually 'confluence') and follow the link.

6. Click the link **Begin migration**. You will know the migration has been successful if you see this reported:

   ```
   1.Migrating users and groups ...
   2.Users and groups migrated successfully!
   ```

   If you encounter errors, please create a support ticket at http://support.atlassian.com and attach your application server logs.

7. Stop Confluence.

8. Start up Confluence and check that you can log in using the admin account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat from the point where you may have gone wrong.

9. Download hibernate_cache_atlassian-user.xml, rename it to `atlassian-user.xml` then copy to your `<INSTALL>/confluence/WEB-INF/classes` directory. It should overwrite the previous `atlassian-user.xml`.

10. Restart Confluence. Check that your users can still log in.

RELATED TOPICS

Understanding User Management in Confluence
Confluence 2.7 Upgrade Guide

Removing a Group

To remove a group,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left panel. A list of all existing groups is displayed along with links to remove them.
3. Click '**Remove**' beside the group you want to remove. You will need to confirm your action before the group is deleted.

RELATED TOPICS

📄 Number of Ancestors to Show in Breadcrumbs

📄 Configuring Number Formats

📄 Configuring HTTP Timeout Settings

Configuring Indexing Language

Configuring Time and Date Formats

Configuring Character Encoding

Thumbnail Settings

Configuring Attachment Size

Recognised System Properties

Removing a User

It is not possible to remove a user if the user is responsible for content on the site (another words, if a user has contributed content. For example, edit, create, or commented within Confluence), because Confluence will need the user information to maintain a history of pages. If you wish to prevent such a user from accessing Confluence, you can deactivate a user so that they can no longer log in to Confluence. Deactivating a user will not remove the content created by them from the site.

You can remove a user if the user has not added or edited any content on the site.

To deactivate or remove a user,

1. First, go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.

   - Select '**Manage Users**' in the left-hand panel.
   - Find the user by searching on the username, full name or email address. You can also click '**Show all users**' to browse the list of all users.
   - Click the username.

1. Now you should be able to see the user's current details and links allowing you to edit them.

2. Click '**Remove**' if the user is not responsible for any content on the site.

3. If the user is responsible for content on the site, you will need to deactivate the user (also known as disabling the user):
   - Remove the user from all groups that have the global '**Can Use**' permission. Click '**Edit Groups**' to remove the person from the group(s).
   - If the specific user has the global '**Can Use**' permission, you will also need to remove this permission.

Number of users and your license
The License Details page tells you how many users your Confluence instance is licensed to support, and how many are currently registered The number of registered users includes only users who have the 'Can Use' global permission. Deactivated users, as described above, are not included.

Deactivating users in Confluence 2.0.x and earlier
Earlier Confluence releases provide a '**Deactivate**' link. If you click this link, the user will be prevented from logging in to the Confluence site. This function was removed from Confluence 2.1.x and later for performance reasons.

Screenshot: Adding and removing users to/from groups

RELATED TOPICS

Removing a Group

Viewing members of a group

Global Groups Overview

Adding or Removing Users in Groups

Enabling or Disabling Public Signup

Searching For and Managing Users

Editing User Details

Changing Usernames

Setting up Anonymous Access

Adding a Group

Security Overview

Troubleshooting LDAP User Management

Adding a New User

Global Permissions Overview

Disabling the Built-In User Management

Removing a User

Administrators Guide Home   Confluence Documentation Home

Setting up Anonymous Access

You can enable anonymous access (also known as public access) to your site by granting the 'Use Confluence' permission to 'Anonymous' users from the 'Administration Console'.

This user category has been created for convenient administration of users who have not logged into the site. Permissions assigned to this group apply to all anonymous users of the site.

To enable public access to your site,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the left panel.
3. Click '**Edit Permissions**'.
4. In the 'Anonymous Access' section, select the '**can use**' check box to enable basic public access to the content on your site.
5. If you selected the '**can use**' check box in the previous step and want to allow public access to user profile views, select the check box in the '**View User Profiles**' section.

   ⚠ You cannot grant the 'View User Profiles' permission independently of the 'Use Confluence' permission.
6. Click the '**Save All**' button to save your changes.

ⓘ Once you grant this permission, further permissions can be granted from the space administration screens to control the viewing and editing privileges of anonymous users. See Space Permissions Overview.

⚠ To disable public access to your site, clear the **can use** check box. If you also have the '**View User Profiles**' check box selected, it must be cleared before saving the configuration changes.

RELATED TOPICS

📄 Adding or Removing Users in Groups

📄 Searching For and Managing Users

📄 Editing User Details

📄 Setting up Anonymous Access

📄 Adding a New User

📄 Global Permissions Overview

📄 Spam Prevention via Captcha

📄 Removing a User

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Viewing members of a group

To view the members of a group,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Manage Groups**' in the left-hand panel. This will list all the existing groups on the site.
3. Click a group name to display all the users in the group.

RELATED TOPICS

📄 Removing a Group

📄 Viewing members of a group

📄 Global Groups Overview

📄 Adding or Removing Users in Groups

📄 Searching For and Managing Users

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

How to Improve User Search Performance

If your Confluence instance contains thousands of user accounts and you are experiencing performance issues when searching for users, the following migration guide is for you.

Background

In Confluence 2.1, we introduced a new system for user management inside Confluence (atlassian-user) that was more powerful than the previous system (OSUser). However, to avoid potential upgrade issues, we continued to use OSUser when storing users in the local Confluence database.

The native atlassian-user storage format provides much more efficient searching, and greatly improves the performance of user administration and Confluence's 'user picker' pop-up. We plan on migrating all Confluence instances to the new format around version 2.6 or 2.7, but until then Confluence instances with large numbers of users can still take advantage of these performance improvements by performing the migration manually.

Migration procedure

> ⛔ Do not use this procedure if you have LDAP user management enabled.

> ℹ️ This guide assumes that you are using Confluence's local users and groups. If you have already configured Confluence for LDAP user/group management and are experiencing user management slowness, please follow the guide for Requesting External User Management Support.

> ✅ **Manual migration no longer required from Confluence 2.7.0**
> This page describes how to perform a manual migration of your users from OSuser to AtlassianUser. For **Confluence 2.7.0 and later**, there is no need to perform a manual migration of your users to the AtlassianUser framework. If you are installing Confluence 2.7.0 or later for the first time, you will automatically receive the AtlassianUser framework. If you are upgrading from an earlier version to Confluence 2.7.0 or later and have not changed the default user management configuration, your users will be automatically migrated. Refer to the details in the Confluence 2.7 Upgrade Guide.

For details of the procedure, refer to Migrating to new User Management.

Restoring Passwords To Recover Admin User Rights

Use this document if you are unable to login as administrator, to manually replace administrator passwords or give users administration rights.

New Confluence User Management

From Confluence 2.7 onwards the user management is handled by AtlassianUser. Hence in the database, Confluence will refer to 'USERS' table to store and refer to its users. When you imported your backup on upgrade, what should happen is the users in the 'OS_USER' table should get copied into 'USERS' table.

If you are still using OSUser please refer to our older document.

Learn more about the algorithm Confluence is using.

Stage One - Identify Administrator

To find out which usernames have admin privileges, connect to your database using a database admin tool such as DBVisualiser. Please download a database admin tool now if you do not have one installed already. Once installed, connect to your database and retrieve the list of administrator usernames with:

```
1. select name from users u, local_members l, groups g where g.groupname = 'confluence-administrators'
   and g.id=l.groupid and u.id=l.userid;
```

Stage Two - Replace Administrator Password

Confluence does not store passwords in plain text in the database, but uses hashes computed from the original password. You instead cut and a paste a hash, rather than the plain password, over the existing password. Below is the hash for the password `admin`

```
1. x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A==
```

To change the password to `admin` for a given username:

1. Shutdown Confluence
2. Connect to your database.
3. The SQL to run is:

```
1.update users set password =
2.'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A=='
3.where name='<USERNAME>';
```

4. Start Confluence
5. Login with your modified username and use password `admin`

For the evaluation embedded database

1. Shut down Confluence.
2. Open <confluence-home>/database/confluence-db.script. Search for:

```
1.INSERT INTO USERS VALUES(
```

3. Replace the password for the appropriate user, you can copy and paste the hash value above.
4. Save the file, and restart.

If No Local Users Exist

In rare circumstances, when local users are deleted and only LDAP users exist, it may be required to insert a user. Here's how to do that:

```
01.insert into users (id, name, password, email, created, fullname) values (1212121, 'admin',
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A==',
'a@a.com', '2009-11-26 17:42:08', 'admin');
02.
03.// Then find out the ID's of Groups
04.SELECT * FROM groups;
05.SELECT * FROM users;
06.
07.// Add group memberships into local_users
08.insert into local_members (userid, groupid) values (<from select above for
user>,<from_select_above_for_conf_users_group>);
09.insert into local_members (userid, groupid) values (<from select above for
user>,<from_select_above_for_conf_admin_group>);
```

Integrating with Crowd

Atlassian's Crowd identity management system can be integrated with Confluence. Please refer to the Crowd documentation on Integrating Crowd with Confluence.

JIRA User Management

- Delegate user management to use JIRA logins
  - Revert from JIRA to internal user management

Additional Information

- JIRA User Management FAQs

Delegate user management to use JIRA logins

If you already have a significant user base set up inside JIRA, it makes sense to connect Confluence to JIRA so that user management is centralised and not duplicated. This document outlines how to delegate Confluence's user authentication and group management to JIRA so that you can use your JIRA users to log in to Confluence.

> ⚠ **Known Issues**
> Before attempting the steps below, please refer to the known issues listed in the troubleshooting section below.

On this page:

- Read Before Proceeding
- Technical Overview
- Step One: Installing Confluence
- Step Two: Setting up a Datasource to JIRA's Database
- Step Three: Installing the JDBC Driver
- Step Four: Modifying osuser.xml
- Step Five: Customising osuser.xml

Read Before Proceeding

1. The examples used in this document are based on **Tomcat Application Server** and the **MySQL database**. The same concepts (but not the verbatim examples) can be applied to other application servers or databases.

2. If JIRA is using LDAP for authentication, you should not use JIRA for Confluence user management. Use Add LDAP Integration instead.

3. Always install Confluence with a new database. Do **not** attempt to use the existing JIRA database, with either JDBC or data source.

4. If you have existing users or groups in Confluence, these users will not be available once you switch to using JIRA's user management. Any existing content will no longer be associated with valid users. Do not add any spaces or content once Confluence installation is complete, apart from verifying that your Confluence instance is up and working, until you have completed the procedure on this page. Users in Confluence will no longer be valid once you switch over to using your JIRA users.

5. If you run into a problem, check the Troubleshooting section below.

Technical Overview

In the configuration described below, Confluence will use JIRA's database for its user and group information. The Confluence application will have two database connections:

1. A connection to the primary database, set up during Confluence installation. This database stores all the normal Confluence data: spaces, pages, comments, etc.
2. A read-only data source connection to JIRA's database, set up after Confluence is installed. Confluence reads information about users and groups from this database.

The reason this works is because both JIRA and Confluence use the same user management library, OSUser. The OSUser database schema is the same in JIRA and Confluence, so Confluence can easily read from JIRA's tables to get the user and group information.

Step One: Installing Confluence

Skip this step if you have already installed Confluence and completed the Setup Wizard.

1. Install Confluence and run the Setup Wizard:
    - If you are running *JIRA standalone* please follow these instructions for installing Confluence.
    - If you have JIRA deployed under your own Tomcat server, please follow these instructions.

2. Ensure that Confluence is running and has been set up, that is, you have completed the Confluence Setup Wizard and verified that you can create pages.

3. Shut down Confluence.

Step Two: Setting up a Datasource to JIRA's Database

In order to delegate all user authentication attempts and group membership queries to JIRA, Confluence needs to be aware of JIRA's database, and hence the user tables in JIRA's database.

In Tomcat this is achieved by specifying JIRA's database as a resource. You will need to declare it inside the `<context>` descriptor you set up in *Step One*.

Notes:

- If there is an existing block of `<Resource>` in the `<context>` descriptor, please do not replace it. Rather, just add the following `<Resource>` block inside the `<context>` descriptor.
- If you are running Confluence EAR/WAR distribution separate to JIRA, or under JIRA standalone 3.3 and later, your Confluence context will be in the `confluence.xml` file.
  If you are running Confluence standalone separate to JIRA, or Confluence WAR/WAR distribution under an older version of JIRA, your Confluence context will be in the `server.xml` file. You should never have a Confluence context in both.
- If you are running Confluence standalone (or Confluence inside a JIRA standalone) and are not sure which version of Tomcat you are using, check your log files. You'll see *INFO: Starting Servlet Engine: Apache Tomcat/5.5.nn* if you are using Tomcat 5.5.
- The DataSource configuration below uses MySQL as an example. You will need to modify these settings according to the database that you are using.
- You should add the appropriate validation check for the connection pool to avoid intermittent problems authenticating.

Sample context descriptor for Tomcat 4.x and 5.0.x

```
01.<Context path="/confluence" docBase="C:/programs/confluence" swallowOutput="true">
02.          <Resource name="jdbc/JiraDS" auth="Container" type="javax.sql.DataSource"/>
03.          <ResourceParams name="jdbc/JiraDS">
04.          <parameter>
05.          <name>username</name>
06.                 <value>your_db_username</value>
07.          </parameter>
08.          <parameter>
09.          <name>password</name>
10.                 <value>your_db_password</value>
11.          </parameter>
12.          <parameter>
13.          <name>driverClassName</name>
14.                 <value>com.mysql.jdbc.Driver</value>
15.          </parameter>
16.          <parameter>
17.          <name>url</name>
18.                 <value>jdbc:mysql://your.domain.com/jira_database_name?autoReconnect=true&lt
;/value>
19.          </parameter>
20.          <parameter>
21.                 <name>factory</name>
22.                 <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
23.          </parameter>
24.          </ResourceParams>
25.</Context>
```

Tomcat 5.5.x

This version of Tomcat has a new syntax for specifying resources. Note that you don't add a new context to `server.xml`, just add the Resource to your existing Context:

```
1.<Context path="/confluence" docBase="C:/programs/confluence" swallowOutput="true">
2.     <Resource name="jdbc/JiraDS" auth="Container" type="javax.sql.DataSource"
3.            username="your_db_username"
4.            password="your_db_password"
5.            driverClassName="com.mysql.jdbc.Driver"
6.            url="jdbc:mysql://your.domain.com/jira_database_name?autoReconnect=true"/>
7.</Context>
```

⚠ Some customers have encountered periodic login failures after delegating user management to JIRA. This usually occurs after a network failure or after the database server reboots. If you do experience such problems after completing this procedure, please refer to Surviving Connection Closures for further information on resolving them. As explained in this document, resolving these issues typically involves adding a `validationQuery` parameter to your JIRA datasource connection definition (such as the example `jdbc/JiraDS` definition above).

Step Three: Installing the JDBC Driver

Ensure that your JDBC driver is on the classpath of your application server. In this example, a jar for the mysql driver should be in the `.../common/lib` folder (or potentially `.../lib` for Tomcat version 6 and beyond).

1. Download the mysql driver from here.

2. Copy the jar file into the `.../common/lib` folder (or `.../lib`).

Step Four: Modifying osuser.xml

⚠ **Please make sure you have completed the Confluence Setup Wizard before performing this step.**

1. Find the `osuser.xml` file in the `/confluence/WEB-INF/classes` folder and open it in a text editor. Comment out the following block of code:

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
03.                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
04.          </provider>
05.          <provider class="bucket.user.providers.CachingAccessProvider">
06.                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
07.                    <property name="chain.configuration.provider.class"
>bucketuser.BucketHibernateConfigProvider</property>
08.          </provider>
09.          <provider class="bucket.user.providers.CachingProfileProvider">
10.                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
11.                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
12.          </provider>
```

2. Uncomment this block:

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
03.                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
04.          </provider>
05.          <provider class="bucket.user.providers.CachingAccessProvider">
06.                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
07.                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
08.          </provider>
09.          <provider class="bucket.user.providers.CachingProfileProvider">
10.                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
11.                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
12.                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
13.          </provider>
```

Your `osuser.xml` should now look like this:

```
01.<opensymphony-user>
02.         <!--
03.                   Authenticators can take properties just like providers.
04.
05.                   This smart authenticator should work for 'most' cases - it dynamically
looks up
06.                   the most appropriate authenticator for the current server.
07.         -->
08.         <authenticator class="com.opensymphony.user.authenticator.SmartAuthenticator"/>
09.
10.         <!-- JIRA User management (with caching) -->
11.         <!-- Note: Do not add any line breaks or spaces when specifying the
chain.classname, otherwise a ClassNotFoundException will be thrown -->
12.
13.         <provider class="bucket.user.providers.CachingCredentialsProvider">
14.                   <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
15.                   <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
16.         </provider>
17.         <provider class="bucket.user.providers.CachingAccessProvider">
18.                   <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
19.                   <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
20.         </provider>
21.         <provider class="bucket.user.providers.CachingProfileProvider">
22.                   <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
23.                   <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
24.                   <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
25.         </provider>
26.
27.         <!--
28.         <provider class="bucket.user.providers.CachingCredentialsProvider">
29.                   <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
30.                   <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
31.         </provider>
32.         <provider class="bucket.user.providers.CachingAccessProvider">
33.                   <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
34.                   <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
35.         </provider>
36.         <provider class="bucket.user.providers.CachingProfileProvider">
37.                   <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
38.                   <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
39.         </provider>
40.         -->
41.
42.         <!--<provider class=
"com.opensymphony.user.provider.memory.MemoryCredentialsProvider" />
43.         <provider class="com.opensymphony.user.provider.memory.MemoryAccessProvider" />
44.         <provider class="com.opensymphony.user.provider.memory.MemoryProfileProvider"
/>-->
45.</opensymphony-user>
```

In this example, **JiraDS** is the name of the JIRA datasource you are sharing with Confluence. If you have changed the name in *Step Two* of this documentation, you will need change all occurrences of the value here too.

You can also download the already configured file here.

Step Five: Customising osuser.xml

In some cases you may need to customise the behaviour of the JiraJdbc classes. You can do this by setting properties within the `osuser.xml` file.

This process is documented here.

Step Six: Modifying atlassian-user.xml

⚠ This step is **only applicable for Confluence 2.7 and later** .

Please comment out or remove the following line from your
`<Confluence-Install>/confluence/WEB-INF/classes/atlassian-user.xml` file:

```
<hibernate name="Hibernate Repository" key="hibernateRepository"  description="Hibernate
Repository" cache="true"/>
```

and add this line instead:

```
<osuser name="OSUser Repository" key="osuserRepository"/>
```

Step Seven: Creating Confluence Groups in JIRA

1. Add the `confluence-users` and `confluence-administrators` groups in JIRA.

2. Add yourself to both these groups.

3. To give your existing JIRA users access to Confluence, you have two options.
    - Option 1: Manually edit the groups of these users inside JIRA and give them membership to one or both of these confluence groups.
    - Option 2: Start up Confluence. Log in using your JIRA account and go to **Administration** and then **Global Permissions**. Now assign the '**can use**' permission to your desired JIRA groups.

    > ⚠️ In order to use Confluence, users must be a member of the `confluence-users` group (or have Confluence 'can use' permission).

Step Eight: Activating External User Management

Since user management is now conducted in JIRA and outside of Confluence, you will need to switch external user management on.

> ⚠️ **Activating external user management will remove user and group management options from Confluence.**
> Your users will also no longer be able to edit their **full name** or **email address** inside Confluence. (If they want to, they would have to do so in JIRA).

To switch external user management on:

1. Log into Confluence using your JIRA account.

2. Go to the **Administration Console** and click **General Configuration** in the left-hand panel

3. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.

4. Select '**ON**' beside '**External User Management**'.

For troubleshooting, see the JIRA Integration FAQ.

RELATED TOPICS

Revert from JIRA to internal user management

Delegate user management to use JIRA logins

Migrating users from Confluence to JIRA

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Revert from JIRA to internal user management

> ℹ️ Check out Crowd for a fully featured user management solution.

Administrators can revert a Confluence instance that uses JIRA for user management back to internal user management. With few users, it is easier to manually recreate the JIRA users and groups in Confluence. For more users, migrate JIRA users and groups into the Confluence database instead.

Option A - Manually Recreate Users In Confluence

This option is too time consuming for hundreds or thousands of users. After completing the reversion, links to users who created or updated Confluence content may go to error screens.

To manually recreate the users, you must first have an instance of Confluence with internal user management and your data.

- If you have made limited customisations to Confluence and migrating would be desirable - follow the upgrade guide and import your data to a new installation.
- Alternatively, if you have made extensive customisations or do not wish to migrate - go to delegating user management to JIRA and remove your JIRA user management by undoing the instructions in revese order. These steps are specific to your instance so cannot be covered here.

Then manually create JIRA's groups and users in Confluence. If you have assigned permissions in Confluence to a group which exists in JIRA, you must create a group in Confluence with the same name. If a user who exists in JIRA has created content or has had permissions assigned to them, you must also create that user in Confluence.

Option B - Transfer JIRA Users & Groups To Confluence

This option manually migrates JIRA users into the Confluence database, but requires knowledge of SQL.

> ⚠️ **Users not using mySQL**
> Users of non-mySQL databases must be experienced enough to modify the SQL to work in their database as examples are provided for mySQL only. If you adapt the SQL to another database, please consider posting the SQL you used to the comments.

> ⚠️ **Users of Confluence 2.0 or older**
> Pre-Confluence 2.0 users may need to modify the instructions to your older schema, or upgrade Confluence. For example, on Confluence 2.1.5 and older, SQL references to the property table must be updated to be called OS_PROPERTYENTRY in all upper-case.

Stage One - Create Backups

Creating backups is the only way to restore your data if something goes wrong.

1. From Confluence, create a full XML backup including attachments.

2. Stop Confluence.

3. Take a backup copy of the Confluence home and install directories.

4. Repeat the above steps for JIRA.

5. From your mySQL admin tool, create a database backup for the JIRA and Confluence databases.

Stage Two - Replace Confluence User Management

Replace the Confluence user and group permissions with JIRA by transferring table content. The SQL provided is specific to mySQL and must be modifed for other databases. For each SQL statement, do a find and replace on the JIRA and Confluence table names to your table names. In the examples, they are called confluence224 and jira364.

1. Login to a DBA tool that can execute SQL on your DB.

2. Erase user and group content from the Confluence DB:

```
1. delete from confluence224.os_propertyentry where entity_name='OSUser_user';
2. delete from confluence224.os_user_group;
3. delete from confluence224.os_group;
4. delete from confluence224.os_user;
```

3. Copy JIRA's groupbase table into Confluence's os_group table:

```
1. insert into confluence224.os_group (id, groupname)
2. select *
3. from jira364.groupbase;
```

4. Copy JIRA's userbase table into Confluence's os_user table:

```
1. insert into confluence224.os_user (id, username, passwd)
2. select *
3. from jira364.userbase;
```

5. Copy JIRA's membershipbase table into Confluence's os_user_group table.

```
1. insert into confluence224.os_user_group (group_id, user_id)
2. select distinct groupbase.id as "group_id", userbase.id as "user_id"
3. from jira364.groupbase, jira364.membershipbase, jira364.userbase
4. where membershipbase.user_name = userbase.username and membershipbase.group_name =
   groupbase.groupname;
```

6. Merge relevant content from JIRA's propertyentry and propertystring tables into Confluence's os_propertyentry table. Some versions of SQL use "0" instead of "false" for boolean values.

```
1. insert into confluence224.os_propertyentry (entity_name, entity_id, entity_key, key_type,
   boolean_val, double_val, string_val, text_val, long_val, int_val, date_val)
2. select 'OSUser_user', propertyentry.entity_id, propertyentry.property_key, 5, false, 0,
   propertystring.propertyvalue, '', 0, 0, null
3. from jira364.propertyentry, jira364.propertystring
4. where propertyentry.entity_name='OSUser' and propertyentry.id=propertystring.id;
```

Stage Three - Revert To Local Management

- If you have made limited customisations to Confluence and migrating would be desirable - Install a new instance of Confluence using the upgrade guide.
- Alternatively, if you have made extensive customisations or do not wish to migrate - go to Delegating User Management to JIRA and remove your JIRA user management by undoing the instructions in revese order. These steps are specific to your instance so cannot be covered here.
- If you wish to migrate to LDAP user management at this point, you can follow the instructions to Migrate to LDAP User Management From OsUser.

Done! Note that the original administrator may not display their groups correctly, however their groups are still present.

LDAP User Management

Overview

Confluence integrates with LDAP user repositories in a variety of ways. Start with the Overview of external user management to learn more, then choose your preferred LDAP connectivity.

There are three choices for LDAP integration:

LDAP Configuration	Internal Users	Internal Groups	LDAP Authentication	LDAP Users	LDAP Groups
LDAP for Authentication Only	✔	✔	✔	✔	✖
LDAP for Group and User Management	✔	✔	✔	✔	✔
LDAP Authentication with OSUser (not supported after 2.7)	✔	✔	✔	✖	✖

In addition to these three options, you may also choose to delegate user management to JIRA or use Crowd for a full-featured solution for multiple applications, directories, and single signon. JIRA user management is a good solution if you already manage users in JIRA.

Additional Information

- Help with troubleshooting external user management
- LDAP FAQ

Check the listing of children pages below for other useful resources.

Add LDAP Integration

> ✔ **Try Atlassian Crowd for powerful LDAP integration**
> Atlassian's Crowd is a web-based single sign-on (SSO) tool that simplifies application provisioning and identity management.

Confluence can delegate user authentication to LDAP and use LDAP group memberships to set the user's Confluence access permissions. This also allows Active Directory (AD) integration. This guide is for both users enabling LDAP, and those upgrading their LDAP scheme to support group management. It applies to LDAP over HTTP and SSL/HTTPS.

Once the LDAP is enabled and LDAP users are using Confluence, you cannot revert to local user management without those users being disabled. However, you can create new local users while using LDAP integration.

Who is this guide for?

If you are using local user management in a version prior to Confluence 2.7, or os_user with authentication-only or jira user management, follow the guide to Migrate to LDAP User Management From OsUser. Otherwise, this is the correct guide for you.

> 🚫 **Integrate only after completing Setup**
> If you are doing an LDAP integration as part of a new install, do not integrate until after you complete the initial setup. You can add LDAP integration after you create the admin user for your instance.

Step 1 - Upgrade Confluence

Please check that you are running the latest version of Confluence. If not, we strongly recommend that you consider upgrading Confluence according to this guide. Confirm that you have upgraded successfully before trying to add LDAP to the new version.

Step 2 - Contact your LDAP/AD Administrator

Integration can only be setup by an administrator confident with running user queries against their LDAP directory. You should request assistance from your LDAP or Active Directory administrator for the following steps.

Step 3 - Check your LDAP server

Confirm this information about your LDAP server.

1. Check your server LDAP version. Supported versions are v2 and v3. Supported LDAP servers include OpenLDAP, Microsoft Active Directory, Novell eDirectory, and any server that uses Java JNDI-LDAP mapping.

2. Your LDAP or Active Directory server must support static groups. This means that the *user DNs* must be stored against a membership attribute inside an LDAP groups. An example of a static group is shown below:

```
01.Dn: CN=Sales and Marketing,CN=Users,DC=ad,DC=atlassian,DC=com
02.objectClass: top; group;
03.cn: Sales and Marketing;
04.distinguishedName: CN=Sales and Marketing,CN=Users,DC=ad,DC=atlassian,DC=com;
05.name: Sales and Marketing;
06....
07.member: CN=John Smith,CN=Users,DC=ad,DC=atlassian,DC=com
08.member: CN=Sally Smith,CN=Users,DC=ad,DC=atlassian,DC=com
09....
```

The membership attribute in this case is **member**, but this is not required. Note that the full DNs of *John and Sally Smith* are listed. If the values against **member** are not full DNs, but are just usernames, then you need to add the flag

```
1.<useUnqualifiedUsernameForMembershipComparison>true
</useUnqualifiedUsernameForMembershipComparison>
```

to your LDAP tag in atlassian-user.xml. Open Directory on OS X uses this configuration.

3. You must not have LDAP groups called 'confluence-users' or 'confluence-administrators'.
4. You must have at least one existing Confluence administrator with System Administrator permissions, whose username does not exist in the LDAP server (see Step 4).

Step 4 - Check the System Administrator account

This step assumes that you have at least one Confluence user account which has System Administrator permissions for your Confluence site. For this account, please check that there isn't an account on your LDAP system that has the *exact same username*.

If there *is* an LDAP account with the exact same username, and you do not have another local Confluence account that has System Administrator permissions rights, then you should perform *one* of the following:

- create another account, that *doesn't* exist on LDAP, to act as the administrator
  OR:
- rename your local Confluence administrator account to use another username that doesn't exist in LDAP
  OR:
- rename your LDAP account

This will ensure that you will have an account *that has sufficient rights* to administer your site after you migrate your users.

Step 5 - Configure your LDAP repository

1. Follow Customising atlassian-user.xml
2. Start up Confluence and check that you can log in using the System Administrator account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat where necessary.
3. If you can't successfully log in with this account, please check that the **username** of this account does not **already exist** in your LDAP server. If *usernames are the same, Confluence recognises LDAP accounts over local Confluence accounts.*
4. If you were using OS user previously, run the user migration. After the migration has run, remove the os user tag from atlassian-user.xml and restart Confluence.

Step 6 - Grant access to LDAP users and groups

To grant Confluence login access to your LDAP groups and users,

1. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the left panel.
3. Click to **Edit Permissions for Groups**.
4. In the textbox to '**Grant Browse Permission**', enter the name of an LDAP group that should have Confluence access. Click '**Add**'.
5. Tick the Can Use box for the LDAP group. If the group is not found, it was not present in your LDAP server.
6. For other LDAP groups that need access to Confluence, add them using the same method.
7. If you are integrating LDAP with Confluence for authentication only, no LDAP groups will appear in Confluence. All the individual LDAP users will have to be manually added to an internal Confluence group with <u>Can Use</u> permissions enabled before they can have access to Confluence.
8. Set up your Confluence page and space permissions for these LDAP groups and users.

Installation complete!

Related Pages

- Confluence LDAP Documentation Index

Troubleshooting

Local user management not retained

If you run into this problem, you may be experiencing this bug.

Check your Confluence version

This documentation applies to the latest version of Confluence. There are a couple of key bugs that have been resolved in Confluence 2.6 or 2.6.1, but that pertain to 2.5.6 and 2.5.7.

1. http://jira.atlassian.com/browse/CONF-9434 relates to hibernate cache=true;
   The xml file supplied here has the hibernate cache set to "true".
2. http://jira.atlassian.com/browse/CONF-9195 relates to the migration step.
   Version 2.6.1 corrects this problem.

More information

- Browse the LDAP FAQ.
- If LDAP users or groups are not displayed in Confluence, try the External User Test tool.
- Check the list of known, unresolved LDAP bugs
- See the comments on this page, from other users who may have left some useful information.
- The 'External User Management' setting in the Confluence Administration Console should be set to OFF. This setting is for using JIRA or Crowd for External User Management.

Support

Failing all else, lodge a support request. Be sure to attach your `atlassian-user.xml`, a copy of the output from the External User Test tool, and a zip of your Confluence logs.

Automatically Add LDAP users to the confluence-users Group

This documentation guides you to use an Authenticator that adds LDAP users to `confluence-users` as they log in.

To use this Authenticator, change the line:

```
1.<authenticator class="com.atlassian.confluence.user.ConfluenceAuthenticator"/>
```

in **<Confluence-Install directory>/confluence/WEB-INF/classes/seraph-config.xml** to:

```
1.<authenticator class="com.atlassian.confluence.user.ConfluenceGroupJoiningAuthenticator"/>
```

As of Confluence 2.2 and later, there is no need to download the class files attached to this page. The change in seraph-config.xml is sufficient.

> ⚠️ If experiencing performance problems when logging in, it may be due to CONF-13754. You can add the files ConfluenceGroupJoiningAuthenticator.class and ConfluenceGroupJoiningAuthenticator$1.class file and put it in `<confluence-installation-directory>/confluence/WEB-INF/classes/com/atlassian/confluence/user`. This will override the built-in version of ConfluenceGroupJoiningAuthenticator in Confluence 3.0.1 and earlier versions of Confluence.

Customising atlassian-user.xml

The LDAP server connection is specified by manually editing the file `atlassian-user.xml`. Confluence 2.3 onwards supports multiple LDAP servers by repeating the instructions below for each server so that there are multiple repositories defined.

Stage 1 - Determine if you are migrating from os_user or atlassian_user

Check your current `.../confluence/WEB-INF/classes/atlassian-user.xml` file. If your os user tag is active, you're using os user in some fashion. Make sure to run the user migration after doing your LDAP details, as described in Migrating to new User Management.

Stage 2 - Configure Connection Details

> ℹ️ To make life easier, use Paddle and Apache Directory Studio to test your LDAP connections without restarting Confluence.

1. Edit the file `.../confluence/WEB-INF/classes/atlassian-user.xml` and configure the connection for either AD or LDAP.
   - Connections in Active Directory
   - Connections in other LDAP servers

2. If your Active Directory Server allows anonymous searches, then you *do not* need to specify a securityPrincipal and securityCredential at all. For an example of how you would configure Confluence to allow anonymous authentication, see Enable Anonymous Authentication in LDAP or Active Directory

3. To connect to LDAP over SSL, see Connect to LDAP via SSL

4. Check your configuration against the example connection details shown below.

   > ⚠️ If you change your ldap key to a different name, you will need to change the cache name to pick it up. This is described in CONFKB181536872.

   ```
   01.<ldap key="ldapRepository" name="LDAP Repository@hecate.atlassian.com" cache="true">
   02.          <host>hecate.atlassian.com</host>
   03.          <port>389</port>
   04.          <securityPrincipal>cn=admin,dc=atlassian,dc=private</securityPrincipal>
   05.          <securityCredential>secret</securityCredential>
   06.          <securityProtocol>plain</securityProtocol>
   07.          <securityAuthentication>simple</securityAuthentication>
   08.          <baseContext>dc=atlassian,dc=private</baseContext>
   09.
   10.....
   ```

Stage 3 - Map LDAP Data Tree

1. Configuring the mappings in `atlassian-user.xml` for either AD or LDAP.
   - Mapping Active Directory
   - Mapping other LDAP servers

2. Check your configuration against the example connection details shown below.

   ```
   01....
   02.
   03.          <baseUserNamespace>dc=staff,dc=perftest,dc=atlassian,dc=private</
   baseUserNamespace>
   04.          <baseGroupNamespace>dc=groups,dc=perftest,dc=atlassian,dc=private</
   baseGroupNamespace>
   05.          <usernameAttribute>cn</usernameAttribute>
   06.          <userSearchFilter>(objectClass=inetorgperson)</userSearchFilter>
   07.          <firstnameAttribute>givenname</firstnameAttribute>
   08.          <surnameAttribute>sn</surnameAttribute>
   09.          <emailAttribute>mail</emailAttribute>
   10.          <groupnameAttribute>cn</groupnameAttribute>
   11.          <groupSearchFilter>(objectClass=groupOfNames)</groupSearchFilter>
   12.          <membershipAttribute>member</membershipAttribute>
   13.</ldap>
   ```

Stage 4 - Directory Search Depth Settings

These are the default settings:

```
1. ...
2. <userSearchAllDepths>false</userSearchAllDepths>
3. <groupSearchAllDepths>false</groupSearchAllDepths>
```

The above settings configure the search depth on users and groups. If you set either attribute to:

- false - Confluence will search only for users/groups directly defined in <baseUserNamespace> and <baseGroupNamespace>.
- true - Confluence will search for users/groups defined in the above namespaces and also in namespaces nested within them. For example, if your users are distributed across multiple namespaces, you should set this option to **true**.

Setting the value to true may have a high cost in performance for large directories, because Confluence will search the whole tree and not just the immediate namespace.

Stage 5 - Optional LDAP Settings

The following settings are the default values for all the options under <ldap> in the `atlassian-user.xml` file. Some of them do not appear in the file normally, but can be added if you need to customise them:

XML tag	Default value	Description
LDAP connection properties		
`securityProtocol`	plain ssl	Allow encrypted (SSL) connections. Can be omitted if anonymous access is available.
`authentication`	simple	Plain-text password transmission. Can be 'none' if anonymous access to the LDAP server is available.
`initialContextFactory`	com.sun.jndi.ldap.LdapCtxFactory	Class name of LDAP provider (default: Sun JNDI)
`batchSize`	100	Size of pages in search results
`poolingOn`	true	Use connection pooling
`connectTimeout`	30000	Timeout in milliseconds when opening new server connections. Default: 30 seconds.
`readTimeout`	60000	Timeout in milliseconds for search and other read operations. Default: 60 seconds.
LDAP connection pool properties		
`initSize`	1	Initial size of connection pool, e.g. number of connections to open at start-up. [1]
`prefSize`	10	Preferred size of connection pool. [1]
`maxSize`	0	Maximum size of connection pool. Zero means no maximum size. [1]
`timeout`	300000	Idle time in milliseconds for a connection before it is removed from the pool. Default: 5 minutes. [1]
`debugLevel`	none	Debug level for logging. [1]
`poolAuthentication`	simple	Authentication for pool connections. [1]
LDAP search properties		
`timeToLive`	0	Time limit on searches in milliseconds. Zero means no limit. [2]
`userSearchAllDepths`	false	Whether user searches should search through the LDAP tree or only for direct children of the DN specified by the userSearchFilter. [3]

groupSearchAllDepths	false	Whether group searches should search through the LDAP tree or only for direct children of the DN specified by the groupSearchFilter. [3]
useUnqualifiedUsernameForMembershipComparison	false	If set to true, Confluence will use the value of the usernameAttribute on the user to check for group membership comparisons instead of the complete distinguished name.

Notes

1. The connection pool properties provided by Atlassian-User correspond with the connection pooling properties in JNDI. See this documentation for further information.
2. More information on this time limit is available on Sun's JNDI tutorial.
3. The "searchAllDepths" parameters toggle between SearchControls.SUBTREE_SCOPE (true) and SearchControls.ONELEVEL_SCOPE (false). See Sun's JNDI tutorial on scope configuration for more information.

To override the default values listed above, you can add the value inside your <ldap> tag after the rest of your configuration:

```
1.<ldap ...>
2....
3.          <groupnameAttribute>cn</groupnameAttribute>
4.          <groupSearchFilter>(objectClass=groupOfNames)</groupSearchFilter>
5.          <membershipAttribute>member</membershipAttribute>
6.          <initSize>20</initSize>      <!-- bigger initial connection pool! -->
7.</ldap>
```

Stage 6 - Optional: Configure LDAP for User Authentication Only

If you'd like to configure LDAP for user management and authentication only, start by configuring LDAP integration with groups then follow the instructions for authentication-only LDAP to disable the group management functionality.

Optional: Configuring Multiple LDAP Repositories

For some LDAP servers, it might be necessary to configure Confluence to connect to multiple LDAP servers. This functionality is available in Confluence 2.3 and above, and has a separate guide: Configuring multiple LDAP repositories.

RELATED TOPICS

Add LDAP Integration
Configuring multiple LDAP repositories
atlassian-user.xml reference
How to write a LDAP search filter, in the *Atlassian Development* space.

Migrate to LDAP User Management From OsUser

Is this guide for me?

Use Add LDAP Integration instead if:

- You are setting up Confluence for the first time.
- You do not need to retain group membership for existing users.
- You are using hibernate user management. To find out, check
  `<confluence-home>/confluence/WEB-INF/classes/atlassian-user.xml`. If there is a hibernate tag but not an os_user tag, you're using hibernate user management.

Background Information

Confluence has three types of user management: os_user (deprecated in confluence 2.7 and later), hibernate (also known as atlassian-user), and LDAP integration.

These correspond to the three tags in atlassian-user.xml:

- <ldap key=...>
- <hibernate name=...>
- <osuser key=...>

In Confluence 2.7, during the upgrade task, users are migrated from os_user to hibernate when Confluence is first started. This migration can also be run manually.

When to Run the User Migration

Under most conditions, you'll want to migrate from OsUser to Hibernate User Management. However, there is one condition under which

you'll want to hold off on doing the automatic user management migration when upgrading.

As described in this jira issue, local groups are not retained when migrating from hibernate user management to LDAP. However, if you are using OsUser for user management and want to switch to LDAP, you can retain group membership by following this procedure:

> ⚠️ If you want to upgrade to 2.7 or later and do LDAP integration, we suggest you do your LDAP integration first, then run the upgrade. If you want to run the upgrade first, you can prevent the user migration from occurring by copying your <confluence-home>/confluence/WEB-INF/classes/atlassian-user.xml file and your osuser.xml files into place before starting Confluence. To check whether the migration has run, look in the 'users' table on the database.

> ℹ️ Make sure your usernames match between os_users and LDAP.

1. Create a test environment. This will create an exact replica of your current version. Use your current version. This will ensure that the automatic user migration does not occur when you restart.
2. Download ldap_hibernate_osuser_atlassian-user.xml, rename it to `atlassian-user.xml` then copy to your `<INSTALL>/confluence/WEB-INF/classes` directory. It should overwrite the previous `atlassian-user.xml`. Your `atlassian-user.xml` file should have LDAP, hibernate and osuser repositories (in this order) enabled.

3. Follow Customising atlassian-user.xml
4. Restart Confluence. Login as an Administrator, and go to this URL:

   ```
   1.<BASEURL>/admin/osuser2atluser.jsp
   ```

   Replace `<BASEURL>` with the URL you currently use to access Confluence. For example, http://confluence.atlassian.com or http://foobar.com/confluence.

   > ✅ The comment "Once the LDAP repository is configured, this migration will ignore users who have the same username as an LDAP user. This will ensure users are not duplicated in Confluence when you have both LDAP and local Confluence users enabled" is meant to describe avoiding CONF-8098. Your users, if they match in LDAP, will be migrated to the LDAP user repository on the database. See the Testing section below to confirm your results.

5. Click the link **Begin migration**. You will know the migration has been successful if you see this reported:

   ```
   1.Migrating users ... Users migrated successfully!
   2.Migrating propertyset data ... Propertyset data migrated successfully!
   3.Migrating groups ... Groups migrated successfully\!
   ```

   If you encounter errors, please create a support ticket at http://support.atlassian.com and attach your application server logs.

6. Stop Confluence.
7. Edit `atlassian-user.xml` file and comment out the `<osuser>` repository.
   Change this line:

   ```
   1.<osuser key="osuserRepository" name="OSUser Repository"/>
   ```

   to this:

   ```
   1.<!-- <osuser key="osuserRepository" name="OSUser Repository"/> -->
   ```

8. Start up Confluence and check that you can login using the admin account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat from there.

Grant access to LDAP users and groups

To grant Confluence login access to your LDAP groups and users:

1. From Confluence, go to Administration > Global Permissions
2. Click to Edit Permissions for Groups
3. In the textbox to Grant Browse Permission, enter the name of an LDAP group that should have Confluence access. Click Add.
4. Tick the Can Use box for the LDAP group. If the group is not found, it was not present in your LDAP server.
5. For other LDAP groups that need access to Confluence, add them using the same method.
6. If you are integrating LDAP with Confluence for authentication only, no LDAP groups will appear in Confluence. All the individual LDAP users will have to be manually added to an internal Confluence group having with Can Use permissions enabled before they can have access to Confluence.
7. Setup your Confluence page and space permissions for these LDAP groups and users.

To setup all LDAP users as members of particular Confluence internal groups, use the LDAP Dynamic Groups Plugin.

Testing

1. Check that groups are associated by visiting a user from the User Browser and logging in.
2. Check your external_entites and external_members tables on your database. If done correctly, these tables should contain the users

who were migrated from os_users and matched in LDAP. You should be able to see their group associations in the external_members table.

```
1.SELECT * FROM external_entities;
2.SELECT * FROM external_members;
```

3. Check your users table. These are the users who were *not* matched in LDAP. It might include the 'admin' user that you made when you originally created the wiki, and perhaps others who've signed up who are not in LDAP. If you need to migrate users from this table into LDAP, check the utility attached to CONF-10654.

Related Pages

- Confluence LDAP Documentation Index
- Add LDAP Integration For User Authentication Only

More information

- LDAP FAQ
- If LDAP users or groups are not displayed in Confluence, download the Paddle diagnostic tool
- List of known, unresolved LDAP bugs
- Comments on this page.

Support

Failing all else, lodge a support request. Be sure to attach your atlassian-user.xml, Paddle logs and a zip of your Confluence logs.

Add LDAP Integration For User Authentication Only

Explanation

If you'd like to configure LDAP for user management and authentication only, start by configuring LDAP integration with groups, and continue to this point to remove group management.

Applies For

- Enabling LDAP for the first time
- Upgrading existing LDAP without enabling group management

Important Points

- LDAP users will be mapped to Confluence. If the Confluence username coincides with the LDAP username, the password lookup is done against LDAP. Group management will happen in Confluence.
- Each LDAP user must be added to an internal Confluence group having Can Use permissions in order for those LDAP users to access Confluence. If a password is created for an LDAP user in Confluence, it will be ignored as the LDAP password will override it.

Instructions

If you do not wish Confluence to retrieve any of your LDAP groups and display them inside Confluence then you can do this by specifying a dummy value for the groupSearchFilter filter in your atlassian-user.xml file. That is, update your atlassian-user.xml file with the following:

```
1.<groupSearchFilter>(objectClass=dummyValue)</groupSearchFilter>
```

An example atlassian-user.xml file:

```
01.<baseUserNamespace>cn=users,dc=ad,dc=atlassian,dc=com</baseUserNamespace>
02.<baseGroupNamespace>ou=groups,dc=ad,dc=atlassian,dc=com</baseGroupNamespace>
03.<usernameAttribute>sAMAccountName</usernameAttribute>
04.<userSearchFilter>(objectClass=user)</userSearchFilter>
05.<firstnameAttribute>givenname</firstnameAttribute>
06.<surnameAttribute>sn</surnameAttribute>
07.<emailAttribute>mail</emailAttribute>
08.<groupnameAttribute>cn</groupnameAttribute>
09.<groupSearchFilter>(objectClass=dummyValue)</groupSearchFilter>
10.<membershipAttribute>member</membershipAttribute>
11.<userSearchAllDepths>false</userSearchAllDepths>
12.<groupSearchAllDepths>false</groupSearchAllDepths>
```

Please note: You will still have to provide a *valid* LDAP DN for `baseGroupNamespace`. Confluence still performs a search for groups quite frequently, so you should use a DN without many child nodes, like an individual user DN.

atlassian-user.xml reference

This page describes the function of each of the tags in an atlassian-user.xml file. These can be added as child tags of the `<ldap>` tag in your atlassian-user.xml file to configure each option.

Developer note: this information is derived from atlassian-user-defaults.xml, which can be found in the source of Atlassian-User under `src/main/resources/`. It also ships in atlassian-user.jar.

Core settings

These settings do not have a default value and must be provided to configure an LDAP connection.

XML tag	Description
LDAP connection properties	
`host`	The host name of the machine running the LDAP server. This must resolve from the machine running Confluence.
`port`	The port number that the LDAP server is running on. This should usually be 389.
`securityPrincipal`	The distinguished name (DN) of a user who is allowed to browse the entire LDAP repository. This can be omitted if the repository has anonymous access enabled.
`securityCredential`	The password for the user configured as the securityPrincipal. This can be omitted if the repository has anonymous access enabled. [1]
`baseContext`	The DN of the top of the LDAP tree that contains both users and groups.
LDAP user mapping properties	
`baseUserNamespace`	The DN at the top of the LDAP tree which contains users. For example: `ou=users,dc=example,dc=com`.
`userSearchFilter`	An LDAP search filter which matches only users under the baseUserNamespace. For example: `(objectClass=person)`. [1]
`usernameAttribute`	The attribute on a user in LDAP which contains the Confluence username. It must be unique across all users. For example: `cn` (OpenLDAP), `sAMAccountName` (AD).
`firstnameAttribute`	The attribute on a user in LDAP which contains the first name of the user. For example: `givenName`.
`surnameAttribute`	The attribute on a user in LDAP which contains the last name of the user. For example: `sn`.
`emailAttribute`	The attribute on a user in LDAP which contains the email address of the user. For example: `mail`.
LDAP group mapping properties	
`baseGroupNamespace`	The DN at the top of the LDAP tree which contains groups. For example: `ou=groups,dc=example,dc=com`.
`groupSearchFilter`	An LDAP search filter which matches only group entities under the baseGroupNamespace. For example: `(objectClass=group)`. [1]
`groupnameAttribute`	The attribute on a group in LDAP which contains the Confluence group name. It must be unique across all groups. For example: `cn`.
`membershipAttribute`	The attribute on a group in LDAP which contains the DN of each member in the group. [2] For example: `member`.

Notes

1. If these values contain ampersands, they must be escaped in the XML file. This is a common situation with LDAP search filters and passwords containing ampersands. For example, the LDAP search filter `(&(objectClass=user)(mail=*@example.com))` would be put in the XML as:
   `<userSearchFilter>(&(objectClass=user)(mail=*@example.com))</userSearchFilter>`.
2. This can also contain the user name of each member of the group. See the `useUnqualifiedUsernameForMembershipComparison` optional configuration setting below.

Optional settings

XML tag	Default value	Description
LDAP connection properties		
`securityProtocol`	plain ssl	Allow encrypted (SSL) connections. Can be omitted if anonymous access is available.
`authentication`	simple	Plain-text password transmission. Can be 'none' if anonymous access to the LDAP server is available.
`initialContextFactory`	com.sun.jndi.ldap.LdapCtxFactory	Class name of LDAP provider (default: Sun JNDI)

batchSize	100	Size of pages in search results
poolingOn	true	Use connection pooling
connectTimeout	30000	Timeout in milliseconds when opening new server connections. Default: 30 seconds.
readTimeout	60000	Timeout in milliseconds for search and other read operations. Default: 60 seconds.
LDAP connection pool properties		
initSize	1	Initial size of connection pool, e.g. number of connections to open at start-up. [1]
prefSize	10	Preferred size of connection pool. [1]
maxSize	0	Maximum size of connection pool. Zero means no maximum size. [1]
timeout	300000	Idle time in milliseconds for a connection before it is removed from the pool. Default: 5 minutes. [1]
debugLevel	none	Debug level for logging. [1]
poolAuthentication	simple	Authentication for pool connections. [1]
LDAP search properties		
timeToLive	0	Time limit on searches in milliseconds. Zero means no limit. [2]
userSearchAllDepths	false	Whether user searches should search through the LDAP tree or only for direct children of the DN specified by the userSearchFilter. [3]
groupSearchAllDepths	false	Whether group searches should search through the LDAP tree or only for direct children of the DN specified by the groupSearchFilter. [3]
useUnqualifiedUsernameForMembershipComparison	false	If set to true, Confluence will use the value of the usernameAttribute on the user to check for group membership comparisons instead of the complete distinguished name.

Notes

1. The connection pool properties provided by Atlassian-User correspond with the connection pooling properties in JNDI. See this documentation for further information.
2. More information on this time limit is available on Sun's JNDI tutorial.
3. The "searchAllDepths" parameters toggle between SearchControls.SUBTREE_SCOPE (true) and SearchControls.ONELEVEL_SCOPE (false). See Sun's JNDI tutorial on scope configuration for more information.

RELATED TOPICS

Customising atlassian-user.xml
Add LDAP Integration
LDAP User Management

Changes in osuser.xml from 1.0.3a to 1.1.x

If you have setup Confluence to:

- delegate user management to JIRA (as per our documentation here) or
- use LDAP/Active directcory authentication (as per our documentation here)

Then you will need to make the following changes to your osuser.xml file:

1. replace

```
1.com.atlassian.confluence.user.providers.Caching
```

with

```
1.bucket.user.providers.Caching
```

Doing this will effectively convert the following:

```
1.com.atlassian.confluence.user.providers.CachingCredentialsProvider
2.com.atlassian.confluence.user.providers.CachingAccessProvider
3.com.atlassian.confluence.user.providers.CachingProfileProvider
```

to

```
1.bucket.user.providers.CachingCredentialsProvider
2.bucket.user.providers.CachingAccessProvider
3.bucket.user.providers.CachingProfileProvider
```

2. and replace

```
1.com.atlassian.confluence.user.ConfluenceHibernateConfigProvider
```

with

```
1.bucket.user.BucketHibernateConfigProvider
```

Alternatively

You can just reconfigure the new osuser.xml with your changes.

Configuring multiple LDAP repositories

In this document::

- Prerequisites
- Configuration
- Side effects
- Cache configuration
- Two connections to the same server
- Single Sign-On Alternative
- Related pages

Prerequisites

Before reading this, ensure you understand the configuration details outlined in Customising atlassian-user.xml. This describes how to configure a single LDAP repository in Confluence, and is prerequisite knowledge for following the instructions below.

For brevity, all examples on this page are partial examples. A complete `atlassian-user.xml` LDAP configuration can be found in Customising atlassian-user.xml.

You will need at least Confluence 2.3 or higher to be able to use these instructions.

Configuration

To configure multiple LDAP repositories in Confluence, put multiple `<ldap>...</ldap>` entries into `confluence/WEB-INF/classes/atlassian-user.xml`.

The order of the entries in the file will be the order that the repositories are searched for users. That is, if a user tries to log in with the username *jsmith*, the first repository in `atlassian-user.xml` will be searched for the user with the username *jsmith*. If no user is found in that repository, the second repository specified in `atlassian-user.xml` will be searched.

Here is a partial configuration that connects Confluence to two different LDAP servers. They are given the identifiers *ldap1* and *ldap2*, and connect to the servers `ldap-sf.example.org` and `ldap-nyc.example.org` respectively.

> ⚠ If you change your ldap key to a different name, you will need to change the cache name to pick it up. This is described in CONFKB181536872.

```
01.<atlassian-user>
02.          <repositories>
03.                    <ldap key="ldap1" name="San Francisco Example Repository" cache="true">
04.                              <host>ldap-sf.example.org</host>
05.                              <port>389</port>
06.
07.                              <!-- ... remainder of server configuration ... -->
08.
09.                    </ldap>
10.
11.                    <ldap key="ldap2" name="New York City Example Repository" cache="true">
12.                              <host>ldap-nyc.example.org</host>
13.                              <port>389</port>
14.
15.                              <!-- ... remainder of server configuration ... -->
16.
17.                    </ldap>
18.
19.                    <hibernate key="hibernate" name="Hibernate Repository" description="Hibernate
Repository" />
20.          </repositories>
21.</atlassian-user>
```

Points to note:

- each server must have a unique *key* attribute
- each server must include the full LDAP configuration, including `baseUserNamespace`, `baseGroupNamespace` and so on
- Confluence's internal repository, the `<hibernate>` repository, must be specified last
- you can include more than two LDAP repositories, but please read the *Side effects* section below.

Side effects

The main side effect of configuring multiple LDAP servers is degrading performance. There are many activities in Confluence where user or group information is retrieved:

- logging in
- user/group searches
- permission checks when viewing or editing a page.

Confluence tries to cache as much information as possible from the LDAP queries, but **almost certainly adding multiple LDAP servers will degrade the performance of the application**. This is especially true if any of the LDAP servers are geographically distant from Confluence, where any LDAP query has a significant latency (> 50 ms roundtrip).

Cache configuration

You will need to configure your
`<Confluence-Installation-Directory>/confluence/web-inf/lib/confluence-2.x.x.jar/confluence-coherence-cac`
or `confluence-coherence-cache-config-clustered.xml` file to add LDAP related caches for every additional LDAP repository being added to `atlassian-user.xml`.

To do this, please add the following block of lines to respective cache configuration file, for each additional LDAP repository being configured:

```
01.<!-- second LDAP repository -->
02.
03.                    <cache-mapping>
04.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.users</
cache-name>
05.                            <scheme-name>user</scheme-name>
06.                    </cache-mapping>
07.
08.                    <cache-mapping>
09.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.users_ro</
cache-name>
10.                            <scheme-name>user</scheme-name>
11.                    </cache-mapping>
12.                    <cache-mapping>
13.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.repository</
cache-name>
14.                            <scheme-name>user</scheme-name>
15.                    </cache-mapping>
16.
17.                    <cache-mapping>
18.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups</
cache-name>
19.                            <scheme-name>user</scheme-name>
20.                    </cache-mapping>
21.                    <cache-mapping>
22.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups_hasMembership</
cache-name>
23.                            <scheme-name>user</scheme-name>
24.                    </cache-mapping>
25.                    <cache-mapping>
26.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups_getGroupsForUser</
cache-name>
27.                            <scheme-name>user</scheme-name>
28.                    </cache-mapping>
29.                    <cache-mapping>
30.                            <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.repositories</
cache-name>
31.                            <scheme-name>user</scheme-name>
32.                    </cache-mapping>
33.
34.                    <!-- END second LDAP Repository -->
```

Please replace the example LDAP key above `ldapRepository2`, with your relevant LDAP key, such that it matches the one defined in your `atlassian-user.xml` file.

Two connections to the same server

It also possible, but not usually recommended, for Confluence to connect twice to the same server. When connecting twice to the same server, **you must not have overlapping group or user namespaces in the LDAP tree**.

Here is an partial configuration, retrieving two separate LDAP user branches, but only one LDAP group branch. To configure only a single group branch, the group filter in the second LDAP repository searches for a non-existent value so it will not return any results. (This is generally fast as long as your LDAP server has an index on objectClass for the given tree section.)

```
01.<atlassian-user>
02.          <repositories>
03.                    <ldap key="ldap1" name="Example Repository, SF user tree" cache="true">
04.                              <host>ldap.example.org</host>
05.                              <port>389</port>
06.                              <!-- ... remainder of connection configuration ... -->
07.
08.                              <!-- user search filter -->
09.                              <baseUserNamespace>cn=San Francisco,dc=ldap,dc=example,dc=org</
baseUserNamespace>
10.                              <userSearchFilter>(objectClass=user)</userSearchFilter>
11.
12.                              <!-- ... remainder of user configuration ... -->
13.
14.                              <!-- group search filter -->
15.                              <baseGroupNamespace>cn=Groups,dc=ldap,dc=example,dc=org</
baseGroupNamespace>
16.                              <groupSearchFilter>(objectClass=group)</groupSearchFilter>
17.
18.                              <!-- ... remainder of server configuration ... -->
19.                    </ldap>
20.                    <ldap key="ldap2" name="Example Repository, NYC user tree" cache="true">
21.                              <host>ldap.example.org</host>
22.                              <port>389</port>
23.                              <!-- ... remainder of connection configuration ... -->
24.
25.                              <!-- user search filter -->
26.                              <baseUserNamespace>cn=New York City,dc=ldap,dc=example,dc=org</
baseUserNamespace>
27.                              <userSearchFilter>(objectClass=user)</userSearchFilter>
28.
29.                              <!-- ... remainder of user configuration ... -->
30.
31.                              <!-- group search filter -->
32.                              <baseGroupNamespace>cn=Groups,dc=ldap,dc=example,dc=org</
baseGroupNamespace>
33.                              <groupSearchFilter>(objectClass=nothing)</groupSearchFilter>
34.
35.                              <!-- ... remainder of server configuration ... -->
36.                    </ldap>
37.                    <hibernate key="hibernate" name="Hibernate Repository" description="Hibernate
Repository" />
38.          </repositories>
39.</atlassian-user>
```

Points to note:

- each repository will have its own connection pool, so **Confluence will use twice as many connections to the LDAP server**
- performance will typically be degraded, as discussed in *Side effects* above
- each server must have a unique *key* attribute
- each server must include the full LDAP configuration, including `baseUserNamespace`, `baseGroupNamespace` and so on
- Confluence's internal repository, the `<hibernate>` repository, must be specified last.

Single Sign-On Alternative

Rather than configuring multiple LDAP repositories, you're able to reduce the overhead of having multiple LDAP servers by setting up an SSO solution as an intermediate user manager. Confluence only needs to lookup the SSO tool once, and that tool then looks up both servers on behalf of Confluence. The SSO tool essentially federates your LDAP servers into a single service and also caches the results across all applications that use single sign-on.

The main advantage is that if a user is already logged into any applications that have single sign-on enabled, then their account details will already be cached by the SSO tool and can normally be served from cache. Only the first application request needs to wait for the LDAP response, with subsequent requests from that or other applications able to use the SSO cache until it expires. If your LDAP servers are high-latency, then you can also reduce latency on cache lookups by locating the SSO tool closer to your Confluence server. Atlassian Crowd is an example SSO solution that already integrates with LDAP and all Atlassian tools.

Related pages

Customising atlassian-user.xml
Add LDAP Integration

Connect to LDAP, JIRA or Other Services Via SSL

Confluence and SSL Services

To serve Confluence pages over SSL, refer to Adding SSL for Secure Logins and Page Security. This document describes connections to SSL Services. The most common are LDAP over SSL and Trusted Applications when JIRA is running over SSL.

Importing SSL Certificates

> ℹ The following commands are in reference to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to this document.

Atlassian User LDAP supports connecting to an LDAP server over SSL/HTTPS. The Trusted Communications protocol requires configuration if JIRA is running over SSL.

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorise the LDAP server's certificate; it will be either the one that was used for signing it, or will come from further up in the trust chain, possibly the root certificate. This is often a self-signed certificate, when both ends of the SSL connection are within the same network. Again, the exact alias is not important.

```
keytool -import -alias serverCert -file RootCert.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias serverCert -file RootCert.crt -keystore
$JAVA_HOME/lib/security/cacerts (Linux/Unix/Mac)
```

2. Import your LDAP or JIRA server's public certificate into the JVM Keystore. This is the certificate that the LDAP server will use to set up the SSL encryption, and you can use any alias of your choosing in place of "JIRAorLDAPServer.crt".

```
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt -keystore
$JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

3. Edit the file in your Confluence install directory, `confluence\WEB-INF\classes\atlassian-user.xml`. Change the value of `securityProtocol` from "plain" to "ssl":

```
<securityProtocol>ssl</securityProtocol>
```

Switch the LDAP connection to the SSL port, if it's different from the default LDAP port. If you're using the most common LDAPS port, set

```
<port>636</port>
```

The keytool will ask you for a password. The default password is "changeit" without the quotes.

4. Verify that the certificate has been added successfully by entering the following command:

```
keytool -list -keystore %JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -list -keystore $JAVA_HOME/jre/lib/security/cacerts (Unix/Linux/Mac)
```

5. Ensure that you've updated JAVA_OPTS to specify the path to the keystore, as specified in Connecting to SSL services, before restarting tomcat/Confluence.
There's no need to specify an alias for Confluence to use; on connecting to the LDAP server, it will search through the keystore to find a certificate to match the key being presented by the server.

Troubleshooting

Check Unable to Connect to SSL Services due to PKIX Path Building Failed sun.security.provider.certpath.SunCertPathBuilderException or the more general SSL Troubleshooting articles for troubleshooting information.

RELATED ARTICLES

JIRA Connecting to SSL Services
Confluence Unable to Connect to SSL Services
Configure Web Proxy Support for Confluence
Adding SSL for Secure Logins and Page Security
Troubleshooting SSL

Disabling the Built-In User Management

> ✅ You only need to follow the instructions on this page if you are using **JIRA** for user management. It disables all the group and user management screens in Confluence.
>
> You can choose to enable this setting for **Crowd** or **LDAP** user management, if you are happy to manage users and groups outside Confluence.

Enabling the "external user management" setting will make user and group management screens **read-only** within Confluence.

This will also prevent users from signing up to the site, as well as editing their name, email and password particulars from within Confluence.

ℹ️ You need to have System Administrator permissions in order to perform this function.

To disable management of users and groups within Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:

   * Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**External User Management**'.
5. Click '**Save**'.

> ⚠️ **Confluence installations using Crowd for user and group management:**
> If Crowd's directory permissions are configured so that Confluence cannot update the Crowd directories, then Confluence's external user management setting *must* be turned on. Otherwise, a 'System Error' will occur when Confluence attempts to write data into Crowd. For more information about integrating Crowd with Confluence, please refer to Integrating Crowd with Atlassian Confluence.

RELATED TOPICS

📄 Changes in osuser.xml from 1.0.3a to 1.1.x

📄 Add LDAP Integration For User Authentication Only

📄 Disabling the Built-In User Management

🏠 Administrators Guide Home 🏠 Confluence Documentation Home

Legacy User Management Documentation

This section contains legacy user management documentation.

* LDAP Authentication with OSUser

LDAP Authentication with OSUser

> ⚠️ **Deprecated for Confluence 2.7.0 and later**
> For Confluence 2.7.0 and above, all user management is performed by AtlassianUser as described in Understanding User Management in Confluence. OSUser user management, described in this document, has been deprecated.
>
> For up-to-date instructions on configuring LDAP authentication, see Add LDAP Integration.

Overview

To configure Confluence to connect to LDAP for user management only, you have two options:

1. Use LDAP for users and authentication, not groups. That's described in Add LDAP integration, and includes an optional step to remove group management. This method is supported in Confluence versions after version 2.7.

   or

2. Create and manage users and groups in Confluence, and use LDAP for authentication only. This is the deprecated OSUser management approach, as described in this document. If a username exists in both Confluence and LDAP, the user uses their LDAP password to log in. This method has been deprecated after Confluence version 2.7

Important Points about Option 2 Above

- Only the password lookup is done against LDAP and *only* if the Confluence username is the same as the LDAP username. Users and user profiles are still managed in Confluence. See technical explanation below.
- A Confluence account must be created for each LDAP user, as users do not automatically have access. For an LDAP user to access Confluence, a site administrator will still need to create an account for them. The password in this Confluence account will be ignored as the LDAP password will override it.

Instructions

Step 1: Modify atlassian-user.xml to use OSUser Management

For Confluence 2.7 and above, the default user repository is the hibernate repository (Atlassian User). To revert to OSUser, you will need to put the OSUser repository tag in the top position so it's the primary user management option. Modify /confluence/WEB-INF/classes/atlassian-user.xml to contain this:

```
1.<atlassian-user>
2.<repositories>
3.<osuser key="osuserRepository" name="OSUser Repository"/>
4.<hibernate name="Hibernate Repository" key="hibernateRepository" description="Hibernate
Repository" cache="true"/>
5.</repositories>
6.</atlassian-user>
```

Note: For Confluence version prior to 2.7, if you have delegated your user management to JIRA, LDAP or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:

- <Installation-Directory>/confluence/WEB-INF/classes/osuser.xml.
- <Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml (if you are upgrading from Confluence 2.2 or later).

Step 2: Open the osuser.xml file located in your home directory under WEB-INF/classes

In the osuser.xml file, the **CredentialsProviders** are responsible for authenticating passwords. The default CachingCredentialsProvider looks in the Confluence database. To enable LDAP aunthentication, you will need to add a **LDAPCredentialsProvider**, so that LDAP users can also be authenticated:

Here's what the default osuser.xml contains:

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
03.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
04.</provider>
05.<provider class="bucket.user.providers.CachingAccessProvider">
06.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
07.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
08.</provider>
09.<provider class="bucket.user.providers.CachingProfileProvider">
10.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
11.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
12.</provider>
```

Step 3: Edit the osuser.xml file as shown below

For Confluence version 2.1 and later:

```
01.<provider class="com.atlassian.confluence.user.ConfluenceLDAPCredentialsProvider">
02.<property name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</property>
03.<property name="java.naming.provider.url">ldap://localhost:389</property>
04.<property name="searchBase">dc=atlassian,dc=com</property>
05.<property name="uidSearchName">cn</property>
06.<!--
07.<property name="java.naming.security.principal">cn=Manager,dc=atlassian,dc=com</property>
08.<property name="java.naming.security.credentials">secret</property>
09.<property name="exclusive-access">true</property>
10.-->
11.</provider>
12.
13.<provider class="bucket.user.providers.CachingCredentialsProvider">
14.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
15.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
16.</provider>
17.<provider class="bucket.user.providers.CachingAccessProvider">
18.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
19.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
20.</provider>
21.<provider class="bucket.user.providers.CachingProfileProvider">
22.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
23.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
24.</provider>
```

For older verisons of Confluence

```
01.<provider class="com.opensymphony.user.provider.ldap.LDAPCredentialsProvider">
02.<property name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</property>
03.<property name="java.naming.provider.url">ldap://localhost:389</property>
04.<property name="searchBase">dc=atlassian,dc=com</property>
05.<property name="uidSearchName">cn</property>
06.<!--
07.<property name="java.naming.security.principal">cn=Manager,dc=atlassian,dc=com</property>
08.<property name="java.naming.security.credentials">secret</property>
09.<property name="exclusive-access">true</property>
10.-->
11.</provider>
12.
13.<provider class="bucket.user.providers.CachingCredentialsProvider">
14.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
15.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
16.</provider>
17.<provider class="bucket.user.providers.CachingAccessProvider">
18.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
19.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
20.</provider>
21.<provider class="bucket.user.providers.CachingProfileProvider">
22.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
23.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
24.</provider>
```

- **Update the following properties to suit your LDAP server:**

    - **url** (currently set to **ldap://localhost:389**)
    - **searchBase** (currently set to **dc=atlassian,dc=com**)
    - **uidSearchName** (currently set to **cn**)

- **If your LDAP server is not configured to allow anonymous lookups, you need to:**

    - remove comment tags
    - enter the username; including searchBase (currently set to *cn=Manager,dc=atlassian,dc=com*)
    - enter password (currently set to *secret*)

ℹ The Credentials (password) checking is a separate operation from user-profile lookups. The profile can be loaded from the Confluence database, but the password is looked up from LDAP. Furthermore, multiple credentials providers can be specified (here, LDAP and OSUser), and if one fails, the other will be used. This allows non-LDAP users to log in with their Confluence password.

How this works

It is useful to have a general idea of how this setup works. This section outlines some consequences of this OSUser implementation and provides some help for people experiencing LDAP connection problems.

Only password-checking for LDAP users is done in Confluence

User profiles are still managed in Confluence (by the CachingProfileProvider in osuser.xml). Only the password lookup is performed against LDAP and only if the Confluence username coincides with a LDAP username. This is because Credentials (password) checking is a separate operation to user-profile lookups. The profile can be loaded from the Confluence database, but the password is looked up from LDAP.

Not all LDAP users have Confluence access

Another effect of this implementation is that LDAP users do not automatically have access to Confluence. A Confluence account must be created for each user wishing to use Confluence.

This is because each Confluence user has a set of groups (for example, 'confluence-users') stored in their profile. Without an associated group, that user can do nothing; not even browse Confluence (that is, they lack the 'use' permission).

Thus, for an LDAP user to use Confluence, a Confluence admin must create an account for them and assign them to a group (typically 'confluence-user'). The password in this Confluence account will be ignored, as the LDAP password will override it.

RELATED TOPICS

📄 Changes in osuser.xml from 1.0.3a to 1.1.x

📄 Add LDAP Integration For User Authentication Only

📄 Disabling the Built-In User Management

🏠 Confluence Documentation Home

Troubleshooting LDAP User Management

Confluence supplies an LDAP connectivity tool (also called '**Paddle**') that will test the LDAP settings in your `atlassian-user.xml` file. This will help you to diagnose problems with LDAP user management, such as:

- LDAP server not responding.
- LDAP settings incorrectly configured in atlassian-user.
- Other issues reported from the LDAP queries run by the testing tool.

LDAP is the protocol used for user management by Active Directory and other LDAP directories.

ℹ **This document applies to Confluence 2.8 and later**
From Confluence version 2.8, the LDAP User Test tool is integrated into the Administration Console. In previous versions of Confluence, this tool was available as a separate utility called '**Paddle**'. If you are running an earlier version of Confluence, please refer to the Confluence 2.7 documentation or choose the relevant Confluence version from the list of previous versions on the documentation home page.

On this page:

- Running the LDAP Connectivity Tool via the Administration Console
- Running the LDAP Connectivity Test Tool outside Confluence
    - Steps in Detail
    - Parameters
- Sample Output
    - Output from a Successful Test
    - Output showing Failure to Connect
    - Output showing No LDAP Configuration
        - RELATED TOPICS

Running the LDAP Connectivity Tool via the Administration Console

⚠ **Note:** You will need to restart your Confluence server each time you make changes to your `atlassian-user.xml` file. To avoid this, consider running the tool outside of Confluence.

To test your LDAP connection settings,

1. First verify that your atlassian-user.xml contains at least one LDAP repository with a key of "ldapRepository".

2. Go to the Confluence '**Administration Console**'. To do this:

   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**LDAP Connectivity Test**' in the '**Administration**' section of the left-hand panel.

   The 'LDAP Connectivity Test' will only appear if your atlassian-user.xml contains an ldap repository with a key of "ldapRepository", as described in the preceding step.

4. The results of the test will appear in the main panel of the screen. Refer to the screenshot below for an example, and to the output samples lower down on this page.

Screenshot: External User Test

Running the LDAP Connectivity Test Tool outside Confluence

You do not need to have Confluence running to run this tool. Instead, you can download this jar and perform the following steps.

Steps in Detail

1. Copy the Paddle jar file into a directory where you have permission to create files.
2. Copy your `atlassian-user.xml` file into the same directory. You will find this file at the following location in your Confluence Installation directory: `<Confluence-Installation>/confluence/WEB-INF/classes/atlassian-user.xml`.
3. Run `java -jar paddle-6.jar`.

Parameters

Paddle supports the following parameters:

Name	Example	Purpose
debug	`java -jar paddle-x.x.jar debug`	Prints DEBUG messages to the console as well as paddle.log.
limit	`java -jar paddle-x.x.jar limit=100`	Sets the limit on the number of results returned by user and group queries. Defaults to 10.

Sample Output

Below are some examples of the output from the user test tool. The output is the same, whether you run the tool via the Administration Console or outside Confluence.

Output from a Successful Test

This is an example of a successful run:

```
001.
##############################################################################################
002.
003.LDAP Support Tool version 4.0
004.
005.
##############################################################################################
006.
007.Connected to server successfully
008.----------------------------------------------------------------
009.TEST 1: Search and list 10 users
010.----------------------------------------------------------------
011.
012.User: cn=foobar2092828899,ou=users,dc=example,dc=com
013.Member of:
014.        cn=ldap-confluence-users236420517,ou=groups,dc=example,dc=com
015.        cn=ldap-confluence-users1970486739,ou=groups,dc=example,dc=com
016.
017.User: cn=foobar1810841610,ou=users,dc=example,dc=com
018.Member of:
019.        cn=ldap-confluence-users646381955,ou=groups,dc=example,dc=com
020.        cn=ldap-confluence-users841429091,ou=groups,dc=example,dc=com
021.
022.User: cn=foobar1678090341,ou=users,dc=example,dc=com
023.Member of:
024.        cn=ldap-confluence-users971262324,ou=groups,dc=example,dc=com
025.        cn=ldap-confluence-users1400937719,ou=groups,dc=example,dc=com
026.
027.User: cn=foobar625635825,ou=users,dc=example,dc=com
028.Member of:
029.        cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
030.        cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
031.
032.User: cn=foobar1740817394,ou=users,dc=example,dc=com
033.Member of:
034.        cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
035.        cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
036.
037.User: cn=foobar39774714,ou=users,dc=example,dc=com
038.Member of:
039.        cn=ldap-confluence-users902164367,ou=groups,dc=example,dc=com
040.        cn=ldap-confluence-users93075839,ou=groups,dc=example,dc=com
041.
042.User: cn=foobar1748690196,ou=users,dc=example,dc=com
043.Member of:
044.        cn=ldap-confluence-users1386043681,ou=groups,dc=example,dc=com
045.        cn=ldap-confluence-users1418978137,ou=groups,dc=example,dc=com
046.
047.User: cn=foobar1958046554,ou=users,dc=example,dc=com
048.Member of:
049.        cn=ldap-confluence-users78102241,ou=groups,dc=example,dc=com
050.        cn=ldap-confluence-users220323310,ou=groups,dc=example,dc=com
051.
```

```
052.User: cn=foobar1263319669,ou=users,dc=example,dc=com
053.Member of:
054.          cn=ldap-confluence-users283478949,ou=groups,dc=example,dc=com
055.          cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
056.
057.User: cn=foobar1541805698,ou=users,dc=example,dc=com
058.Member of:
059.          cn=ldap-confluence-users1736591783,ou=groups,dc=example,dc=com
060.          cn=ldap-confluence-users91103896,ou=groups,dc=example,dc=com
061.
062.------------------------------------------------------------------
063.TEST 2: Search and list 10 groups
064.------------------------------------------------------------------
065.
066.Group: cn=ldap-confluence-users16368779,ou=groups,dc=example,dc=com
067.Members:
068.          cn=foobar1588244344,ou=users,dc=example,dc=com
069.          cn=foobar997638232,ou=users,dc=example,dc=com
070.          cn=foobar1474202427,ou=users,dc=example,dc=com
071.          cn=foobar1371091481,ou=users,dc=example,dc=com
072.          cn=foobar794627990,ou=users,dc=example,dc=com
073.          cn=foobar904159761,ou=users,dc=example,dc=com
074.          cn=foobar1992607123,ou=users,dc=example,dc=com
075.          cn=foobar839977082,ou=users,dc=example,dc=com
076.          cn=foobar725629285,ou=users,dc=example,dc=com
077.          cn=foobar1317948366,ou=users,dc=example,dc=com
078.
079.Group: cn=ldap-confluence-users1721354031,ou=groups,dc=example,dc=com
080.Members:
081.          cn=foobar991592891,ou=users,dc=example,dc=com
082.          cn=foobar1710975716,ou=users,dc=example,dc=com
083.          cn=foobar1505673129,ou=users,dc=example,dc=com
084.          cn=foobar1924966176,ou=users,dc=example,dc=com
085.          cn=foobar799226338,ou=users,dc=example,dc=com
086.          cn=foobar1999491022,ou=users,dc=example,dc=com
087.          cn=foobar14075397,ou=users,dc=example,dc=com
088.          cn=foobar539697111,ou=users,dc=example,dc=com
089.          cn=foobar112108820,ou=users,dc=example,dc=com
090.          cn=foobar1369772211,ou=users,dc=example,dc=com
091.
092.Group: cn=ldap-confluence-users1418978137,ou=groups,dc=example,dc=com
093.Members:
094.          cn=foobar1588906497,ou=users,dc=example,dc=com
095.          cn=foobar713515675,ou=users,dc=example,dc=com
096.          cn=foobar111198746,ou=users,dc=example,dc=com
097.          cn=foobar1413209068,ou=users,dc=example,dc=com
098.          cn=foobar1748690196,ou=users,dc=example,dc=com
099.          cn=foobar1650270608,ou=users,dc=example,dc=com
100.          cn=foobar1083182874,ou=users,dc=example,dc=com
101.          cn=foobar397077941,ou=users,dc=example,dc=com
102.          cn=foobar982743351,ou=users,dc=example,dc=com
103.          cn=foobar1305199694,ou=users,dc=example,dc=com
104.
105.Group: cn=ldap-confluence-users1970486739,ou=groups,dc=example,dc=com
106.Members:
107.          cn=foobar547480036,ou=users,dc=example,dc=com
108.          cn=foobar2092828899,ou=users,dc=example,dc=com
109.          cn=foobar1150791207,ou=users,dc=example,dc=com
110.          cn=foobar1648997840,ou=users,dc=example,dc=com
111.          cn=foobar632840294,ou=users,dc=example,dc=com
112.          cn=foobar1688785372,ou=users,dc=example,dc=com
113.          cn=foobar1402211544,ou=users,dc=example,dc=com
114.          cn=foobar1360676914,ou=users,dc=example,dc=com
115.          cn=foobar1900838417,ou=users,dc=example,dc=com
116.          cn=foobar174731773,ou=users,dc=example,dc=com
117.
118.Group: cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
119.Members:
120.          cn=foobar763847686,ou=users,dc=example,dc=com
121.          cn=foobar532641707,ou=users,dc=example,dc=com
122.          cn=foobar1478107215,ou=users,dc=example,dc=com
123.          cn=foobar1457615713,ou=users,dc=example,dc=com
124.          cn=foobar1740817394,ou=users,dc=example,dc=com
125.          cn=foobar477128147,ou=users,dc=example,dc=com
126.          cn=foobar1686683086,ou=users,dc=example,dc=com
127.          cn=foobar625635825,ou=users,dc=example,dc=com
128.          cn=foobar122705808,ou=users,dc=example,dc=com
129.          cn=foobar827130393,ou=users,dc=example,dc=com
130.
131.Group: cn=ldap-confluence-users93075839,ou=groups,dc=example,dc=com
132.Members:
133.          cn=foobar1197670610,ou=users,dc=example,dc=com
```

```
134.          cn=foobar1478107215,ou=users,dc=example,dc=com
135.          cn=foobar477128147,ou=users,dc=example,dc=com
136.          cn=foobar122705808,ou=users,dc=example,dc=com
137.          cn=foobar847443321,ou=users,dc=example,dc=com
138.          cn=foobar39774714,ou=users,dc=example,dc=com
139.          cn=foobar797719296,ou=users,dc=example,dc=com
140.          cn=foobar1695434944,ou=users,dc=example,dc=com
141.          cn=foobar1251231239,ou=users,dc=example,dc=com
142.          cn=foobar2524660,ou=users,dc=example,dc=com
143.
144.Group: cn=ldap-confluence-users91103896,ou=groups,dc=example,dc=com
145.Members:
146.          cn=foobar991592891,ou=users,dc=example,dc=com
147.          cn=foobar1688694241,ou=users,dc=example,dc=com
148.          cn=foobar1924966176,ou=users,dc=example,dc=com
149.          cn=foobar1070601909,ou=users,dc=example,dc=com
150.          cn=foobar95317734,ou=users,dc=example,dc=com
151.          cn=foobar1999491022,ou=users,dc=example,dc=com
152.          cn=foobar62701314,ou=users,dc=example,dc=com
153.          cn=foobar14075397,ou=users,dc=example,dc=com
154.          cn=foobar1541805698,ou=users,dc=example,dc=com
155.          cn=foobar369641998,ou=users,dc=example,dc=com
156.
157.Group: cn=ldap-confluence-users78102241,ou=groups,dc=example,dc=com
158.Members:
159.          cn=foobar997638232,ou=users,dc=example,dc=com
160.          cn=foobar1588244344,ou=users,dc=example,dc=com
161.          cn=foobar1783755629,ou=users,dc=example,dc=com
162.          cn=foobar1371091481,ou=users,dc=example,dc=com
163.          cn=foobar904159761,ou=users,dc=example,dc=com
164.          cn=foobar794627990,ou=users,dc=example,dc=com
165.          cn=foobar104918275,ou=users,dc=example,dc=com
166.          cn=foobar839977082,ou=users,dc=example,dc=com
167.          cn=foobar725629285,ou=users,dc=example,dc=com
168.          cn=foobar1958046554,ou=users,dc=example,dc=com
169.
170.Group: cn=ldap-confluence-users236420517,ou=groups,dc=example,dc=com
171.Members:
172.          cn=foobar799557307,ou=users,dc=example,dc=com
173.          cn=foobar1258451414,ou=users,dc=example,dc=com
174.          cn=foobar1041699029,ou=users,dc=example,dc=com
175.          cn=foobar2092828899,ou=users,dc=example,dc=com
176.          cn=foobar547480036,ou=users,dc=example,dc=com
177.          cn=foobar2111497410,ou=users,dc=example,dc=com
178.          cn=foobar1648997840,ou=users,dc=example,dc=com
179.          cn=foobar1776604828,ou=users,dc=example,dc=com
180.          cn=foobar1121496402,ou=users,dc=example,dc=com
181.          cn=foobar1954769270,ou=users,dc=example,dc=com
182.
183.Group: cn=ldap-confluence-users283478949,ou=groups,dc=example,dc=com
184.Members:
185.          cn=foobar1588906497,ou=users,dc=example,dc=com
186.          cn=foobar111198746,ou=users,dc=example,dc=com
187.          cn=foobar357608135,ou=users,dc=example,dc=com
188.          cn=foobar1413209068,ou=users,dc=example,dc=com
189.          cn=foobar1263319669,ou=users,dc=example,dc=com
190.          cn=foobar1650270608,ou=users,dc=example,dc=com
191.          cn=foobar397077941,ou=users,dc=example,dc=com
192.          cn=foobar2116500618,ou=users,dc=example,dc=com
193.          cn=foobar1668222621,ou=users,dc=example,dc=com
194.          cn=foobar1032804410,ou=users,dc=example,dc=com
```

Output showing Failure to Connect

This is an example of the output when the LDAP server is not available:

```
01.
#######################################################################################
02.
03.LDAP Support Tool version 4.0
04.
05.
#######################################################################################
06.
07.Failed to connect to LDAP server: Failed to connect to LDAP server with provider URL ldap:
//ldap.example.com:389: ldap.example.com:389
08.com.atlassian.paddle.connection.ConnectionException: Failed to connect to LDAP server with
provider URL ldap://ldap.example.com:389: ldap.example.com:389
09.        at
com.atlassian.paddle.connection.DefaultConnectionFactory.createFriendlyLdapException(DefaultConnectio
55)
10.        at
com.atlassian.paddle.connection.DefaultConnectionFactory.createConnection(DefaultConnectionFactory.ja
35)
11.        at com.atlassian.paddle.task.DefaultTaskRunner.runTask(DefaultTaskRunner.java:33)
12.        at com.atlassian.paddle.Paddle.testConnectivity(Paddle.java:98)
13..
14..
15..
16.<<<snip>>>
```

Output showing No LDAP Configuration

This is an example of the output when the `atlassian-user.xml` file is not correctly configured for LDAP user management:

```
01.
#######################################################################################
02.
03.LDAP Support Tool version 4.0
04.
05.
#######################################################################################
06.
07.Error reading atlassian-user.xml file: No LDAP settings found in XML configuration.
08.com.atlassian.paddle.configuration.ConfigurationException: No LDAP settings found in XML
configuration.
09.        at
com.atlassian.paddle.configuration.AtlassianUserConfiguration.setAtlassianUserXml(AtlassianUserConfig
44)
10.        at
com.atlassian.paddle.configuration.AtlassianUserConfiguration.(AtlassianUserConfiguration.java:
21)
11.        at com.atlassian.paddle.Paddle.testConnectivity(Paddle.java:79)
12.        at
com.atlassian.confluence.admin.actions.ExternalUserConnectivityAction.execute(ExternalUserConnectivit
23)
13..
14..
15..
16.<<<snip>>>
```

RELATED TOPICS

Requesting External User Management Support
LDAP User Management
User Management

Troubleshooting the "Not Permitted" Screen under LDAP Integration

If you have set up Confluence with AtlassianUser integration and attempted a login with an LDAP/AD account and got a page titled **"Not Permitted"** here are the steps to troubleshoot this:

- Have you assiged USE permission to the relevant LDAP groups from the Administration > Global Permissons page?

Before an LDAP user can login and use Confluence, the LDAP group they belong to must be given USE permission directly. (Please note: nested groups is not supported in Confluence as yet, so you cannot specify the parent group. If you want this feature, please vote for it here.)

- Does an account exist on your LDAP/AD server that has the same name as your local admin account? (for example, is there an acccount on LDAP called 'admin'?)

If so, then you will not be able to login with your local admin account once you enable LDAP integration. To rectify this, you need to either rename your LDAP admin account or rollback your LDAP integration and create another Confluence admin account.

- You have assigned USE permission to the relevant LDAP groups, but LDAP users in those groups still get "Not Permitted"?

Here, you need to check if Confluence is actually aware that your LDAP users belong to those LDAP groups. To work this out, here is what we need from you:

1. Login as the local admin account you created when you first set up Confluence.
2. Enable profiling by appending `?profile=on` to the end of a Confluence URL (say the URL of the dashboard, if you happen to be on that page) and hit enter to refresh the page. This setting will now be enabled and cause additional information to be written to your log files to help us diagnose the problem.
3. Now browse to the `Administration > Manage Users` screen and do a search for any LDAP user.
4. Now click on the user to view their details (including the groups they belong to).
5. Now submit a support ticket at http://support.atlassian.com and attach:
   - A screenshot of the user details page.
   - Your server logs files.

If you are feeling brave, you can attempt to decipher the logs yourself. Here's how:

1. Your logs should display something similar to this:

```
1.[0ms] -
com.atlassian.user.impl.ldap.adaptor.LDAPStaticGroupAdaptor_search((&(objectClass=groupOfNames
2.    [0ms] - com.atlassian.user.impl.ldap.repository.DefaultLDAPRepository_getLDAPContext
3.    [0ms] -
com.atlassian.user.impl.ldap.adaptor.LDAPStaticGroupAdaptor_search_JNDI_RAW_((&(objectClass=gr
```

   This means that Confluence is using this LDAP search filter **(&(objectClass=groupOfNames)(member=cn=confadmin,ou=users,ou=people,ou=functest,dc=atlassian,dc=com))** to find the groups that the user 'confadmin' belongs to. Obviously, the objectClass and member attributes may differ in your install, but the filter should be similar.

2. Connect to your LDAP/AD server with JXplorer or an LDAP tool of your choice, and issue the above filter and check that you get the results you expect. This should help you to identify if and why the filter is incorrect and what it should be. Please add what you find in this step to the support ticket if you are unable to resolve it from here.

Migrating users from Confluence to JIRA

There is currently no way to delegate user management from JIRA to Confluence. So, if you are in a situation where your users are defined in Confluence and would like to take advantage of Confluence's ability to use JIRA user management, you will need to transfer all of your existing Confluence users into JIRA. You can do this manually, or if you have a large number of users, you can use the attached XML-RPC script.

> ⛔ You should backup your JIRA installation before running this script. This is an experimental script that has not gone through the same extensive testing as the Confluence and JIRA products.

Getting the migration tool

- Download the attached rpc-tools-0.9.zip.
- Extract the ZIP file to a temporary directory on your computer.

Running the migration tool

- Back up your JIRA database.
- If you do not have an 'admin' username with password 'admin' in both JIRA and Confluence, create it now.
- Ensure JIRA and Confluence have remote API access enabled. In both applications, it is configured in the *General Configuration* screen in Administration.
- Edit `connection.properties` in the extracted ZIP file to the set the XML-RPC URLs for JIRA and Confluence. The XML-RPC URLs in the default file correspond to the following application base URLs:
  - JIRA - http://localhost:8080
  - Confluence - http://localhost:8080/confluence
- Run `java -jar tools.jar` in the extracted ZIP file directory.

A successful run will generate output like the following:

```
$ java -jar tools.jar
- Transferring group: confluence-administrators
- Transferring group: confluence-users
- Transferring user: testuser1
- Transferring user: testuser2
```

Things to note:

- This script requires that both the Confluence and JIRA remote APIs are available and accessible to a username 'admin', password 'admin'. You can temporarily add this user to both system to run the script, then remove it afterwards.
- A random password will be assigned to each user that is transferred because it is not possible to access password information via

the XML-RPC API. Therefore they will have retrieve a new password via the password reminder.

- JIRA does not allow users to sign up if they do not have a valid email address. Therefore we will assign them an email address of type username@example.com to any Confluence users that do not have a email address. This will allow you to find the users and help them to create a password or change the email address.

Trouble running the script?

- Ensure you have created a user 'admin' with password 'admin' in both Confluence and JIRA.
- Ensure both applications have remote API access enabled (see above).
- Ensure you have patched the Confluence RPC plugin if running Confluence 2.0.x.

If you're still having trouble, please raise a support request, and include a copy of the error you're getting.

Requesting External User Management Support

- LDAP Troubleshooting Resources
- Problems During Initial Setup
- Complex Authentication or Performance Problems
    - Confluence server
    - Confluence configuration files
    - User management system
    - Using Active Directory for LDAP?
    - Diagnostics

LDAP Troubleshooting Resources

- Add LDAP Integration
- Troubleshooting User Management and Login Issues

If the above resources don't help, continue below.

Problems During Initial Setup

Open a Support Ticket and include:

- Refer to Troubleshooting LDAP User Management. Run the tests and paste the output in your support ticket.
- Download an LDAP browser to make sure you've got the right values. Atlassian recommends LDAP Studio. Include screenshots of your user and group DN's.
- Attach your `atlassian-user.xml` file.

Complex Authentication or Performance Problems

Open a Support Ticket and include:

Confluence server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Global Permissions`, if you are having problems with logging in
- Take a screenshot of the Space permissions page, if you are having problems with space or page permissions.

Confluence configuration files

- Attach a copy of `atlassian-user.xml`, found in confluence/WEB-INF/classes
- Attach a copy of `osuser.xml`, found in confluence/WEB-INF/classes.
- If you have implemented a custom authenticator or in any way modified `seraph-config.xml` or `seraph-paths.xml`, please provide the modified files as well.

User management system

- What is the name and version of your LDAP server?
- Does your LDAP server use dynamic or static groups?

Using Active Directory for LDAP?

Please include LDAP Studio Entry Editor snaphots with the information specified on this page

Diagnostics

- Enable profiling (as described here)
- Enable detailed user management logging by editing `confluence/WEB-INF/classes/log4j.properties`:

Change this section:

```
01.###
02.# Atlassian User
03.###
04.#log4j.logger.com.atlassian.user=DEBUG, confluencelog
05.#log4j.additivity.com.atlassian.user=false
06.#log4j.logger.com.atlassian.confluence.user=DEBUG,console
07.#log4j.additivity.com.atlassian.confluence.user=false
08.#log4j.logger.bucket.user=DEBUG,console
09.#log4j.additivity.bucket.user=false
```

To this:

```
01.###
02.# Atlassian User
03.###
04.log4j.logger.com.atlassian.user=DEBUG, confluencelog
05.log4j.additivity.com.atlassian.user=false
06.log4j.logger.com.atlassian.confluence.user=DEBUG,confluencelog
07.log4j.additivity.com.atlassian.confluence.user=false
08.log4j.logger.bucket.user=DEBUG,confluencelog
09.log4j.additivity.bucket.user=false
10.log4j.logger.com.atlassian.seraph=DEBUG, confluencelog
11.log4j.additivity.com.atlassian.seraph=false
12.log4j.logger.com.opensymphony.user=DEBUG, confluencelog
13.log4j.additivity.com.opensymphony.user=false
```

- After enabling both the above, please attempt a Confluence LDAP account login and attach a copy of the log files that are produced when the problem occurs. To do this, locate your install directory or exploded WAR directory, then zip the full /logs subdirectory into a single file for us to examine. An example location might be confluence-2.2.2-std/logs.

> ⚠ If you are using LDAP, run the External User Test tool and attach a copy of the output to the support ticket.

Paddle

From Confluence version 2.8 the '**External User Test**' tool is integrated into the Administration Console, and you can also run it outside of Confluence. In previous versions of Confluence, this tool was available only as a separate utility called '**Paddle**'.

- For **Confluence 2.8 and later**, please refer to Troubleshooting LDAP User Management.
- If you are running an earlier version of Confluence, please refer to the Confluence 2.7 documentation or choose the relevant Confluence version from the list of previous versions on the documentation home page.
- Sometimes for troubleshooting purposes it may still be useful to use the stand-alone Paddle so you don't have to start and stop Confluence each time you test. Place your atlassian-user.xml file in the same directory as the jar file and run: `java -jar paddle-2.0.jar`.

> ⚠ **Be aware that:**
> The Atlassian-User optional setting `useUnqualifiedUsernameForMembershipComparison` is not supported by Paddle, so the tests may not be correct if this option is being used. In this case, the Connectivity test must can be done using Confluence.

Understanding User Management in Confluence

> ℹ **Looking for help with your user management configuration?** See Requesting External User Management Support.

> ✓ **Try Atlassian Crowd for powerful user administration**
> Atlassian's Crowd is a web-based single sign-on (SSO) tool that simplifies application provisioning and identity management, including LDAP integration. For more information, please see the Crowd documentation on Integrating Crowd with Confluence.

On this page:

- Components of User Management
- Authentication
    - Seraph
    - XML-RPC and SOAP Authentication
    - Password Authentication and User Management
- Confluence User Management Frameworks
    - AtlassianUser

- - OSUser
- Related pages

Components of User Management

1. **Authentication** - determining what user identity is making a request to Confluence.
2. **User management** - storing and retrieving core information about users.
3. **Group membership** - storing and retrieving groups, and group membership.
4. **Profile information** - providing metadata associated with users.

It's important to understand that these are separate components of the user management system. The term *LDAP integration* is not really meaningful, because you could use LDAP repository for any or all of the above tasks.

For example, in OSUser authentication can be performed against different repository to that used for group membership queries. In AtlassianUser, authentication and group membership can be retrieved from LDAP, but profile information is still stored in the Confluence database.

Authentication

Seraph

Almost all authentication in Confluence (and JIRA) is performed through Seraph, Atlasian's open source web authentication framework. The goal of seraph is to provide a simple, extensible authentication system that we can use on any application server.

Seraph is implemented as a servlet filter. Its sole job is, given a web request, to associate that request with a particular user (or no user if the request is anonymous). It supports several methods of authentication, including HTTP Basic Authentication, form-based authentication, and looking up credentials already stored in the user's session.

Seraph performs no user management itself. It merely checks the credentials of the incoming request, and delegates any user-management functions (looking up a user, checking a user's password is correct) to Confluence's user-management subsystem.

If you were looking to integrate Confluence with a Single Sign-On (SSO) infrastructure, you would do so by writing a custom Seraph authenticator (and in fact, many customers have done so).

XML-RPC and SOAP Authentication

Normally, requests for Confluence's remote API will include an authentication token as the first argument. With this method of authentication, XML-RPC and SOAP authentication requests are checked directly against the user-management framework, and tokens are assigned directly by the remote API subsystem. These requests do not pass through Seraph authenticators.

However, if the provided token argument is blank, Seraph will be used as a fallback authentication method for remote API requests. So, to use a custom Seraph authenticator with XML-RPC or SOAP requests, ensure that you pass an empty string as the authentication token to remote API methods.

Password Authentication and User Management

By default, password authentication is delegated from Seraph to the user management system. This is not necessary, however. Single Sign-On systems may have no password authentication at all, and get all the necessary credentials from the SSO provider.

Confluence User Management Frameworks

The rest of this document covers the user management frameworks used by Confluence: **AtlassianUser** and **OSUser**.

AtlassianUser

AtlassianUser is a new user and group management framework developed by Atlassian, and is the core framework used in Confluence since version 2.1. AtlassianUser was developed with the following goals in mind:

- Support LDAP as a fully functional repository for authentication, group management and profile information (profile information not yet implemented).
- Compatibility with both JIRA and Confluence (JIRA support not yet implemented).
- Be simple to configure.

AtlassianUser provides user, group and profile management services to Confluence. It does so by delegating tasks to configurable *repositories*. Multiple repositories can be configured, so for example Confluence can draw user information from both the database *and* an LDAP server.

Default Configuration

Configuration of AtlassianUser is done through the `<<confluence-install>>\confluence\WEB-INF\classes\atlassian-user.xml` file. See the `atlassian-user.xml` reference page. (In Confluence 2.1, configuration of AtlassianUser is done through the `atlassianUserContext.xml` file.)

For Confluence 2.7.0 and later:

- All user management is performed by AtlassianUser's native providers.

- OSUser delegation is still supported for customers who rely on the OSUser/JIRA bridge or the old-style OSUser LDAP support.
- Refer to the Confluence 2.7 Upgrade Guide for details of the automatic migration which may occur during the upgrade process, for customers who are using the standard user management framework.

For Confluence 2.6.x and earlier:

- Confluence's AtlassianUser configuration delegates all user, group, profile and password authentication to OSUser.

Database (Hibernate) Support in AtlassianUser

AtlassianUser can store user, group and profile data directly in Confluence's database. This is the default behaviour for Confluence 2.7.0 and later.

LDAP Support in AtlassianUser

AtlassianUser currently supports *password authentication*, *user management* and *group management* with an LDAP server. Follow the instructions on configuring AtlassianUser LDAP integration.

At this point, only read-only access to LDAP is planned. Java's JNDI-LDAP interface does not support updating an LDAP repository, and the administration tools that come with LDAP servers such as Microsoft Active Directory are generally comprehensive and already available in enterprise IT departments.

JIRA Integration via AtlassianUser

AtlassianUser will not support delegating Confluence user management to JIRA. Instead, our goal is to implement AtlassianUser as the JIRA user management framework as well. Once this is done, both Confluence and JIRA can use the same LDAP server for their authentication and group management.

OSUser

OpenSymphony User was Confluence's core user management framework until it was replaced by AtlassianUser in version 2.1. OSUser is still supported through AtlassianUser's OSUser repositories. OSUser is also built around the model of pluggable providers, but its LDAP support is limited.

OSUser Database (Hibernate) Providers

In its default configuration, Confluence's OSUser providers store a list of users and groups together with profile information in tables in the Confluence database:

- os_user (authentication)
- os_group (group membership)
- os_user_group (group membership)
- os_propertyentry (profile information)

The hashed password in the os_user table is used to authenticate the user unless LDAP support is enabled. The os_user_group table is queried for group membership information.

OSUser configuration is controlled through the `<<confluence-install>>\confluence\WEB-INF\classes\osuser.xml` file.

LDAP Support in OSUser

OSUser only supports *authentication* against an LDAP server. That is, you can check user passwords against LDAP, but all other user information must be shadowed in the Confluence database. Follow the instructions on configuring OSUser LDAP authentication.

If you need support for LDAP user information or group membership as well, you should use AtlassianUser instead (see above).

Delegating User Management to JIRA via OSUser

Confluence can use OSUser to retrieve information for *authentication*, *group membership* and *profile information* from JIRA.

If you look at the discussion of OSUser's implementation above, you can see how this can works pretty easily. A data source to JIRA database is configured in Confluence which lets Confluence read directly from JIRA's os_* tables. For example, when a user is created in JIRA, the username and password goes in the os_user table in the JIRA database. Confluence looks at the same table in the JIRA database to authenticate the user.

Access to the JIRA database is read-only. For this reason, Confluence maintains a subset of the user's *profile information* locally in the Confluence database (things like last login time and user preferences that Confluence needs to be able to modify) .

Follow the instructions on configuring OSUser delegation to JIRA.

Related pages

HTTP authentication with Seraph
Single Sign-on Integration with JIRA and Confluence
Add LDAP Integration
LDAP Authentication with OSUser
Delegate user management to use JIRA logins

Migrating to new User Management
Confluence 2.7 Upgrade Guide
atlassian-user.xml reference

User Management Frequently Asked Questions

This page has been split into the LDAP FAQ and JIRA Integration FAQ.